

IBM System Storage DS Storage Manager Version 10.8

Installation and Host Support Guide



Note

Before using this information and the product it supports, read the information in “Notices” on page 317.

This edition applies to version 10 modification 83 of the IBM DS Storage Manager, and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces GA32-2221-02.

© **Copyright IBM Corporation 2012, 2013.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	vii
----------------	------------

Tables	ix
---------------	-----------

About this document	xi
----------------------------	-----------

What's new in this release	xii
Related documentation	xiii
Storage Manager documentation on the IBM website	xiii
Storage Manager online help and diagnostics	xiv
Finding Storage Manager software, controller firmware, and readme files	xiv
IBM System Storage Productivity Center	xv
Essential websites for support information	xv
Getting information, help, and service	xvi
Before you call	xvi
Using the documentation	xvii
Software service and support	xvii
Hardware service and support	xvii
Notices and statements in this document	xvii
Receiving product updates and support notifications	xviii

Chapter 1. Preparing for installation . . . 1

Introduction	1
Storage Manager software	1
Storage Manager software components	2
Supported controller firmware	2
Types of installation configurations	2
Network configuration	3
Direct-attached and SAN-attached configurations	6
Setting up controller addresses for software installation	7
Setting up IP addresses for storage subsystem controllers	7
Setting up an IP address with the DHCP/BOOTP server	8
Assigning static TCP/IP addresses to a storage subsystem using factory-default management port TCP/IP address	9
Assigning static TCP/IP addresses storage subsystem using an in-band management connection	10
Assigning static TCP/IP addresses using the storage subsystem controller serial port Service Interface	11

Chapter 2. The Storage Manager interface . . . 13

Enterprise Management window	13
Using the Devices tab	15
Using the Setup tab	19
Subsystem Management window	19
Opening the Subsystem Management window	20

Using the Summary tab	21
Using the Storage and Copy Services tab	21
Using the Host Mappings tab	24
Using the Hardware tab	25
Using the Setup tab	26
Managing multiple software versions	26

Chapter 3. Installing Storage Manager 27

Preinstallation requirements	27
Installing the Storage Manager packages automatically with the installation wizard	28
Installing Storage Manager with a console window in Linux, AIX, HP-UX, and Solaris	31
Installing Storage Manager packages manually	32
Software installation sequence	32
Installing Storage Manager manually	32
Uninstalling Storage Manager	33
Uninstalling Storage Manager on a Windows operating system	33
Uninstalling Storage Manager on a Linux, AIX, or Solaris operating system	34
Completing the Storage Manager installation	34
Performing an automatic discovery of storage subsystems	34
Performing a manual discovery of storage subsystems	36
Setting a storage subsystem management password	36
Naming storage subsystems	37
Setting alert notifications	37
Managing iSCSI settings	39
Downloading controller firmware, NVSRAM, ESM firmware	43
Downloading drive firmware	49
Storage Manager premium features	50
Enabling the premium feature trial version	51
Enabling the permanent premium feature	52
Obtaining the premium feature enable identifier	52
Generating the feature key file	53
Enabling the premium feature	53
Disabling premium features	54
Saving the storage subsystem profile	54

Chapter 4. Configuring storage . . . 55

Storage partitioning overview	55
Using the Task Assistant	56
Drives supported by IBM System Storage DS Storage Manager	56
Rules for selecting drives when creating a RAID array	57
Solid state drive (SSD) attributes	59
T10PI capable drive attributes	60
Full Disk Encryption (FDE) attributes	66
Configuring disk storage	67
Creating a disk pool	72

Creating an array	74
Redundant array of independent disks (RAID)	74
Creating a standard logical drive	77
Creating a thin logical drive	78
About Dynamic Capacity Expansion	78
Viewing Operations in Progress	79
Configuring global hot-spare drives	80
Defining a default host operating system	81
Defining a host group	83
Defining heterogeneous hosts	83
Defining the host and host ports	84
Mapping LUNs	84
Mapping LUNs to a new Host or Host Group	84
Adding LUNs to an existing Host or Host Group	85
Configuring the IBM System Storage for IBM i	85
Configuring and using optional premium features	86
About Enhanced FlashCopy	86
About FlashCopy	87
Using VolumeCopy	88
Using Enhanced Remote Mirroring	88
Using Enhanced Global Mirroring	88
Using Performance Read Cache	88
Using Full Disk Encryption	89
Using other features	89
Using controller cache memory	89
Using Persistent Reservations	91
Using Media Scan	91
Tuning storage subsystems	95
Maximizing throughput with load balancing	96
Balancing the Fibre Channel I/O load	97
Optimizing the I/O transfer rate	97
Optimizing the I/O request rate	98
Using the Storage Manager command-line interface and Script Editor	100
Storage Manager command-line interface	100
Using the Script Editor	100

Chapter 5. Configuring hosts 103

Booting a host operating system using SAN boot	103
Overview of multipath drivers	105
Using multipath drivers to automatically manage logical drive fail-over and fail-back	115
Using host bus adapters	118
Installing a multipath driver	120
AIX multipath drivers	129
Linux Device Mapper Multipath driver	129
Linux RDAC (MPP) driver	138
MAC OS multipath driver	143
Veritas DMP driver	143
HP-UX PV-links	143
Solaris failover drivers	148
Identifying devices	159
Using the SMdevices utility	159
Identifying devices on AIX hosts	160
Configuring devices	162
Using the hot_add utility	162
Using the SMrepassist utility	163
Stopping and restarting the host-agent software	163
Setting the queue depth for hdisk devices	164
Disabling cache mirroring	165

Using dynamic capacity expansion and dynamic logical drive expansion	166
Veritas Storage Foundation with SUSE Linux Enterprise Server	167
Veritas Storage Foundation 5.0 with Red Hat Enterprise Linux	167
Checking LUN size	169
Redistributing logical drives	169
Replacing hot-swap HBAs	171
Settings for the Windows DSM and Linux RDAC	181
Configuration Settings for Path Congestion Detection and Online/Offline Path States	186
Setting up details on the DS5000 storage system and AIX host to support T10PI	189
Set up the DS5K storage box	189
Set up the AIX host	189

Chapter 6. Working with full disk encryption. 191

Full disk encryption	192
Securing data against a breach	192
Choosing local or external security key management	193
Using security keys	194
Using secure erase	206
FDE security authorizations	207
FDE terminology	209
Before you begin	210
Installing and configuring the DS TKLM Proxy Code server	210
Starting, stopping, and restarting the DS TKLM Proxy Code server	211
Modifying the DS TKLM Proxy Code server configuration file	212
Installing the DS TKLM Proxy Code	215
Configuring disk encryption with FDE drives	216
Installing FDE drives	217
Enabling premium features	217
Securing a RAID array	224
Unlocking disk drives	229
Migrating storage subsystems (head-swap) with FDE drives	232
Erasing disk drives	235
Global hot-spare disk drives	238
Log files	239
Frequently asked questions	239
Securing arrays	240
Secure erase	240
Local security key management	241
External security key management	241
Premium features	241
Global hot-spare drives	242
Boot support	242
Locked and unlocked states	242
Backup and recovery	242
Other	243

Chapter 7. Troubleshooting 245

Critical event problem solving	245
Retrieve trace buffers	260

Configuration database validation	261
Database save/restore	262
DS Diagnostic Data Capture (DDC)	263
Recovery steps	264
DDC MEL events	265
Resolving disk array errors on AIX	266
IBM DS Storage Manager - Password Reset	270

Appendix A. Host bus adapter settings 271

Adjusting HBA settings	271
Accessing HBA settings through Fast!UTIL	271
Default host bus adapter settings	272
Advanced HBA settings	273
QLogic host bus adapter settings	274
JNI and QLogic host bus adapter settings	279
JNI HBA card settings	279
QLogic HBA settings	283

Appendix B. Using a storage subsystem with a VMware ESX Server configuration 285

Sample configuration	285
Software requirements	286
Management station	286
Host (VMware ESX Server).	286
Hardware requirements	286
VMware ESX Server restrictions	287
Other VMware ESX Server host information	289
Configuring storage subsystems for VMware ESX Server	289
Cross-connect configuration for VMware connections	289
Mapping LUNs to a storage partition on VMware ESX Server	289
Verifying the storage configuration for VMware	290

Appendix C. Using the Storage Manager with high-availability cluster services. 291

General information	291
Using cluster services on AIX systems	291
High-Availability Cluster Multi-Processing	291
Parallel System Support Programs and General Parallel File System	293
GPFS, PSSP, and HACMP cluster configuration diagrams	293
Using cluster services on HP-UX systems	299
Using cluster services on Solaris systems	300
General Solaris requirements	300
System dependencies	300

Appendix D. Viewing and setting AIX Object Data Manager (ODM) attributes 301

Attribute definitions	301
Using the lsattr command to view ODM attributes	306

Appendix E. About VDS/VSS provider 309

Appendix F. Installing SMI-S provider 311

Appendix G. Accessibility 313

Notices 317

Trademarks	319
Important notes	320

Glossary 321

Index 333

Figures

1. Sample network using network-managed and host-agent-managed storage subsystems	3
2. Parts of the Enterprise Management window	15
3. Disk Pool Automatic Configuration	59
4. Protection Information (P) check points	61
5. Enabling T10 PI on a logical drive	63
6. RAID drive - Protection Information (T10 PI) - enabled	64
7. Example - Logical Drive 4 of RAID array 4 - T10PI not enabled	65
8. Disabling T10PI	66
9. FDE capable RAID array - security details	67
10. Assigning a port identifier for IBM i	85
11. Selecting IBM i as the host type.	86
12. The Script Editor window	101
13. I/O flow in an optimal single path	106
14. I/O flow in optimal two path	107
15. Use of one path when the other fails.	108
16. Failover of I/O in a single path environment	109
17. Failover of I/O in a multipath environment	110
18. All paths to controller fail in AVT/ADT and RDAC failover modes.	111
19. All paths to controller fail in ALUA failover mode. First five minutes of the failover.	112
20. All paths to the controller fail in ALUA mode. Five minutes into the failure	113
21. Host HBA to storage subsystem controller multipath sample configuration for all multipath drivers except AIX fcp_array and Solaris RDAC	117
22. Host HBA to storage subsystem controller multipath sample configuration for AIX fcp_array and Solaris RDAC multipath drivers.	118
23. One-to-one zoning scheme	120
24. One-to-two zoning scheme	120
25. Security-enabled FDE drives: With the correct authorizations in place, the reading and writing of data occurs in Unlocked state	196
26. A security-enabled FDE drive is removed from the storage subsystem: Without correct authorizations, a stolen FDE disk cannot be unlocked, and the data remains encrypted	197
27. Changing the security key	199
28. Changing the security key - Complete	200
29. Drive properties - Secure FDE drive	201
30. Select file - LockKeyID	203
31. Drive properties - Unsecured FDE drive	204
32. Secure erase process	207
33. External security key management topology	211
34. Sample VMware ESX Server configuration	285
35. Cross-connect configuration for VMware connections	289
36. Cluster configuration with single storage subsystem—one to four partitions.	294
37. Cluster configuration with three storage subsystems—one partition per subsystem	295
38. Cluster configuration with four storage subsystems—one partition per subsystem	296
39. RVSD cluster configuration with two storage subsystems—two partitions per subsystem.	297
40. HACMP/GPFS cluster configuration with one storage subsystem—one partition	298
41. HACMP/GPFS cluster configuration with two storage subsystems—two partitions per subsystem	299

Tables

1. Data shown in the table view	17	29. Troubleshooting the Device Mapper	137
2. Adding a Storage Subsystem.	18	30. Description of mppUtil parameters	141
3. Removing Storage Subsystems	18	31. Sample record of logical drive preferred and alternate paths	146
4. Simultaneously removing multiple subsystems	19	32. Configuration parameters of the failover driver	182
5. Nodes on the Logical tab	22	33. Parameters for Wait time settings	185
6. Types of nodes in the Topology pane	24	34. Configuration Settings for the Path Congestion Detection	186
7. Node information in the Defined Mappings pane	25	35. Security authorizations	207
8. Installation sequence of Storage Manager software packages	32	36. Full disk encryption terminology	209
9. Examples of Storage Manager package install commands	33	37. Proxy configuration file properties	212
10. Storage Manager package installation verify commands	33	38. Critical events	245
11. Summary of supported drive types, interfaces and capabilities	57	39. Recovery Step 2.	264
12. Protection Information metadata (8 bytes)	60	40. Recovery Step 4.	264
13. Types of drives that can be used in arrays and disk pools	68	41. Recovery Step 5.	265
14. Copy Services support by array and disk pool	68	42. DDC MEL events	265
15. Reserved capacity in a disk pool	70	43. Disk array errors	266
16. List of features supported in an array or a disk pool.	70	44. QLogic model QLA234x, QLA24xx, QLE2462, QLE2460, QLE2560, QLE2562, QMI2572, QMI3572, QMI2582	274
17. RAID level descriptions	75	45. QLogic model QL220x (for BIOS V1.81) host bus adapter settings by operating system	278
18. Maximum supported Performance Read Cache size per installed controller cache	89	46. Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562	279
19. Errors discovered during a media scan	93	47. Configuration settings for FCE-1063/FCE2-1063/FCE-6410/FCE2-6410	281
20. Performance Monitor tuning options in the Subsystem Management window	95	48. Configuration settings for FCI-1063	282
21. Load balancing policies supported by operating systems	96	49. Configuration settings for FC64-1063	283
22. Failover mode for each Operating System	114	50. Configuration settings for QL2342	284
23. Multipath driver by operating system	116	51. Attributes for dar devices	301
24. Number of paths each multipath driver supports by operating system	118	52. Attributes for dac devices	303
25. dsmUtil parameters	126	53. Attributes for hdisk devices.	303
26. Minimum version required for each component	130	54. Example 1: Displaying the attribute settings for a dar	306
27. Attributes and parameter values in the multipath.conf file	135	55. Example 2: Displaying the attribute settings for a dac	306
28. Options and parameters for the multipath command.	137	56. Example 3: Displaying the attribute settings for an hdisk	307
		57. Storage Manager alternate keyboard operations	314

About this document

Lists the tasks for installing the Storage Manager and supporting the host computer. The tasks include determining the hardware and software, integrating hardware with the network, installing the Storage Manager software, and using features of the Storage Manager.

This document provides information about how to plan, install, configure, and work with the IBM® System Storage® DS Storage Manager Version 10.83 or later, and for storage subsystems with controller firmware version 7.8x.xx.xx or later. If the IBM DS Storage Manager is earlier than 10.83, refer to IBM System Storage DS Storage Manager Version 10 - Installation and Host Support Guide. This document is intended for system and storage administrators who are responsible for installing storage administration software. To install and manage a storage subsystem with Storage Manager, you must have knowledge of redundant array of independent disks (RAID), small computer system interface (SCSI), Fibre Channel, and SATA technology. You must also have working knowledge of the applicable operating systems that are used with the management software.

Note: The screenshots in the guide are illustrative and may differ from the actual UI depending on the Storage Manager and controller firmware versions.

Throughout this document, the term *Storage Manager* refers to all host software release levels.

Use this document to perform the following tasks:

- Determine the hardware and software that are required to install Storage Manager.
- Integrate the necessary hardware components into your network.
- Install the Storage Manager software.
- Upgrade controller firmware, if necessary.
- Identify and use storage-management features that are unique to your installation.

Important: Check the Storage Manager readme files for any updates to the list of supported operating systems. See “Finding Storage Manager software, controller firmware, and readme files” on page xiv for more information about how to access the Storage Manager readme files on the web.

For information about terminology, see the Help section of the Storage Manager Enterprise Management window, the Subsystem Management window, or the “Glossary” on page 321.

As you read this document, it is important to understand the distinction between the following two terms:

Management station

A management station is a system that is used to manage the storage subsystem. You can attach it to the storage subsystem in either of the following ways:

- Through a TCP/IP Ethernet connection to the controllers in the storage subsystem

- Through a TCP/IP connection to the host-agent software that is installed on a host computer, which in turn is either directly attached to the storage subsystem through the Fibre Channel I/O path or through a TCP/IP Ethernet connection to the controllers

Host computer

A host computer is a system that is directly attached to the storage subsystem through a Fibre Channel I/O path. This system is used to perform the following tasks:

- Serve data (typically in the form of files) from the storage subsystem
- Function as a connection point to the storage subsystem for a remote-management station

Note:

1. The terms *host* and *host computer* are used interchangeably throughout this document.
2. A host computer can also function as a management station.

What's new in this release

Lists features offered in DS Storage Manager version 10.86.

The features supported in IBM DS Storage Manager version 10.86 are -

Software XOR Engine: Software XOR Engine enhances subsystem performance while running high-bandwidth applications. When a large amount of data is written to the subsystem, the hardware parity service may decrease the write rate. Firmware Parity service augments hardware parity; it increases the amount of data that can be written to the subsystem. When hardware parity is saturated, new parity jobs are routed to Firmware Parity. This service is available on controllers where bandwidth limits must be raised to support target performance levels. RAID 5 parity, and RAID 6 P and Q parities (RAID 5 XOR parity is identical to RAID 6 P parity), can be computed in Firmware Parity. Services available in hardware RAID Parity Assist (RPA) can also be implemented (For example: copy, compare, set). Interaction between Software XOR Engine and Protection Information (PI) is similar to the interaction between Protection Information (PI) and hardware parity service. The PI fields in the parity blocks are computed using the PI fields in the data blocks. A published list of errors related to data alignment is included with the Crystal Beach 3 (CB3) RPA hardware. All source and destination addresses must be on 64-byte boundaries. If data is not correctly aligned, the CB3 chip can get locked up. FPE is used instead of CB3 for RPA requests that are not correctly aligned. The most commonly implemented unalignment is parity computation for cache blocks that include PI. CB3 is also used to compute CRCs for cache backup to flash memory that occurs during a power failure. Similarly, CRCs are verified during cache restore. If a CRC operation is not correctly aligned, the FPE must be used to compute or verify the CRC.

Workload Capture: You can use this mechanism to analyze storage subsystem performance. This analysis is used to tune the host system and storage subsystem for optimum performance under individual site conditions. This release includes functionality for offline analysis of host I/O traffic and performance statistics. The results of the analysis can be used to configure changes in order to increase performance. Future releases will include functionality for capturing and formatting data for ASUP, and configuring changes real-time. Only development and support personnel can control the workload analysis feature and data

collection. It should be used as a non-intrusive process during normal storage system operations. The capabilities and functionality of Workload Capture is similar to other available performance monitoring tools. Workload Capture does not use the SYMBol interface, and does not require the IBM DS Storage Manager running.

View Only Password Management: With the view permission management feature of IBM DS Storage Manager, the existing single security level is extended to two security permission levels: view and configuration change permissions. Authentication of storage subsystem management is ensured at the SYMBol procedure level. SYMBol procedures that modify storage subsystem configuration and perform destructive operations are called 'active' procedures, while SYMBol procedures that report storage subsystem states and configurations are called 'passive procedures'. You must enter the subsystem password to invoke 'active' SYMBol procedures. In this case, this password is defined as 'administrative subsystem password'. FDE requires additional security measures to prevent unauthorized users from modifying or retrieving the FDE lock key. As an additional security measure, you must use the 'administrative subsystem password' to launch the Subsystem Management Window. Users who are not authorized to modify storage subsystem configurations, but are authorized to view storage configurations and monitor health conditions, can perform 'view only' subsystem management operations. The subsystem password for 'view only permission' is defined as 'view subsystem password'. Unlike 'administrative subsystem password', which is enforced at SYMBol procedure request level, the authentication of 'view subsystem password' is managed at the management-session level. The storage subsystem provides a persistent repository for 'view subsystem password' and password validation.

Related documentation

In addition to the information in this document, the resources that are described in the following sections are available.

Storage Manager documentation on the IBM website

Lists the software guides available on the support portal and the procedure to access those.

The following documentation is available for download (PDF) on the IBM website:

- *IBM System Storage DS[®] Storage Manager Command Line Interface and Script Commands Programming Guide*
- *IBM System Storage DS Storage Manager Copy Services User's Guide*
- *IBM System Storage DS4000[®] Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview*

To access these documents and other IBM System Storage documentation from the IBM Support Portal, complete the following steps.

Note: The first time that you access the IBM Support Portal, you must choose the product category, product family, and model numbers for your storage subsystems. The next time you access the IBM Support Portal, the products you selected initially are preloaded by the website, and only the links for your products are displayed. To change or add to your product list, click the **Manage my product lists** link.

1. Go to <http://www.ibm.com/support/entry/portal>.

2. Under **Choose your products**, expand **Hardware**.
3. Click **System Storage > Disk systems > Mid-range disk systems** (for DS4000 or DS5000 storage subsystems) or **Entry-level disk systems** (for DS3000 storage subsystems), and check the box for your storage subsystem.
4. Under **Choose your task**, click **Documentation**.
5. Under **See your results**, click **View your page**.
6. In the **Product documentation** box, click the link for the publication that you want to access.

Storage Manager online help and diagnostics

You can access the help systems from the Enterprise Management and Subsystem Management windows in the Storage Manager by clicking **Help** on the toolbar or pressing F1.

Enterprise Management Help window

Use this online help system to learn more about working with the entire management domain.

Subsystem Management Help window

Use this online help system to learn more about managing individual storage subsystems.

After you install Storage Manager, consider installing the host bus adapter (HBA) management and diagnostic application, if available. The QLogic SANsurfer and Emulex HBAnyware applications are diagnostic programs that you can use to verify the status of the I/O connections before you use the storage subsystem.

If your storage subsystem is connected to a Fibre Channel HBA in the host server in a SAN environment, consider purchasing the IBM Tivoli® Storage Manager software application for SAN management and troubleshooting.

Finding Storage Manager software, controller firmware, and readme files

Lists steps to download the latest versions of the Storage Manager software and controller firmware from the support portal.

The Storage Manager software and controller firmware versions are available on the Storage Manager DVD, and can also be downloaded from the web.

Important: Before you install Storage Manager, review the readme file. Updated readme files contain the latest device-driver versions, firmware levels, limitations, and other information that is not found in this document.

To find firmware and readme files on the IBM Support Portal, complete the following steps:

Note: The first time that you access the IBM Support Portal, you must choose the product category, product family, and model numbers for your storage subsystems. The next time you access the IBM Support Portal, the products you selected initially are preloaded by the website, and only the links for your products are displayed. To change or add to your product list, click the **Manage my product lists** link.

1. Go to <http://www.ibm.com/support/entry/portal>.
2. Under **Choose your products**, expand **Hardware**.

3. Click **System Storage > Disk systems > Mid-range disk systems**, and select the check box for your storage subsystem.
4. Click **Select OS**, check the corresponding box for your operating system, and click **Submit**.
5. Under **Choose your task**, click **Downloads**.
6. Under **See your results**, click **View your page**.
7. In the **Product documentation** box, click the link for the publication that you want to access.

IBM System Storage Productivity Center

The IBM System Storage Productivity Center (SSPC) is an integrated hardware and software solution that provides a single point of entry for managing IBM System Storage DS3000 systems, DS4000 systems, DS5000 systems, DS8000® systems, IBM System Storage SAN Volume Controller clusters, and other components of your data storage infrastructure. Therefore, you can use the IBM System Storage Productivity Center to manage multiple IBM System Storage product configurations from a single management interface.

To learn how to integrate Storage Manager with the IBM System Storage Productivity Center, see the IBM System Storage Productivity Center Information Center at the following website:

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

Essential websites for support information

Lists websites that contain information about the Storage Manager, firmware, and NVSRAM, and downloads.

The most up-to-date information about your IBM storage subsystems and Storage Manager, including documentation and the most recent software, firmware, and NVSRAM downloads, can be found at the following websites:

IBM System Storage Disk Storage Systems

Find links to software and firmware downloads, readme files, and support pages for all IBM System Storage disk storage systems:

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

IBM System Storage Interoperation Center (SSIC)

Find technical support information for your specific storage subsystem and host configuration, including the latest firmware versions for your system, with this interactive web-based utility:

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

IBM DS3000, DS4000, DS5000, and BladeCenter® Boot Disk System Premium Feature Activation

Activate a premium feature with this web-based utility:

<http://www.ibm.com/storage/fasttkeys>

IBM System Storage Productivity Center

Find the latest documentation for the IBM System Storage Productivity Center, a new system that is designed to provide a central management console for IBM System Storage DS3000, DS4000, DS5000, DS8000, and SAN Volume Controller:

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

IBM System Storage Support

Find the latest support information for host operating systems, HBAs, clustering, storage area networks (SANs), Storage Manager software and controller firmware:

www.ibm.com/systems/support/storage

Storage Area Network (SAN) Support

Find information about using SAN switches, including links to SAN documentation:

www.ibm.com/systems/support/storage/san

Support for IBM System p[®] AIX 5L[™] and Linux servers

Find the latest support information for System p AIX[®], Linux, BladeCenter, and i5/OS[™] servers:

www.ibm.com/systems/support/supportsite.wss/brandmain?brandind=5000025

Support for IBM System x[®] servers

Find the latest support information for System x Intel- and AMD-based servers:

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

System p and AIX Information Center

Find information about how to use AIX with System p and POWER[®] servers:

[publib.boulder.ibm.com/infocenter/pseries/index.jsp?](http://publib.boulder.ibm.com/infocenter/pseries/index.jsp)

IBM System Storage products

Find information about all IBM System Storage products:

www.ibm.com/systems/storage

IBM Publications Center

Find IBM publications:

www.ibm.com/shop/publications/order/

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your system, and whom to call for service, if it is necessary.

Before you call

Lists steps to solve the problem yourself before you call IBM support.

Before you call, take these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system.
- Check for technical information, hints, tips, and new device drivers at the IBM System Storage Disk Support website pages that are listed in this section.
- Use an IBM discussion forum on the IBM website to ask questions.

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the Storage Manager online help or in the documents that are provided with your system and software. The information that comes with your system also describes the diagnostic tests that you can perform. Most subsystems, operating systems, and programs come with information that contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the information for the operating system or program.

Using the documentation

Information about your IBM system and preinstalled software, if any, is available in the documents that come with your system; this includes printed books, online documents, readme files, and help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software.

Software service and support

Lists IBM support websites where you can get telephone assistance about software problems.

Through IBM Support Line, for a fee you can get telephone assistance with usage, configuration, and software problems. For information about which products are supported by Support Line in your country or region, go to the following website:

www.ibm.com/services/sl/products

For more information about the IBM Support Line and other IBM services, go to the following websites:

- www.ibm.com/services
- www.ibm.com/planetwide

Hardware service and support

Contains the website for hardware service with its available time in U.S, Canada, and U.K.

You can receive hardware service through IBM Integrated Technology Services or through your IBM reseller, if your reseller is authorized by IBM to provide warranty service. Go to the following website for support telephone numbers:

www.ibm.com/planetwide

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

Notices and statements in this document

This document contains the following notices, which highlight key information:

- **Note:** These notices provide important tips, guidance, or advice.
- **Important:** These notices provide information that might help you avoid inconvenient or problem situations.

- **Attention:** These notices indicate possible damage to programs, devices, or data. An attention notice is placed just before the instruction or situation in which damage could occur.
- **Caution:** These statements indicate situations that can be potentially hazardous to you. A caution statement is placed just before the description of a potentially hazardous procedure step or situation.

Receiving product updates and support notifications

Informs you how to receive product update notifications about the Storage Manager software, controller firmware, ESM firmware, and drive firmware.

Download the latest versions of the following packages at the time of initial installation and when product updates become available:

- Storage Manager host software
- Storage subsystem controller firmware
- Drive storage expansion enclosure ESM firmware
- Drive firmware

Important: Keep your systems current with the latest firmware and other product updates by subscribing to receive support notifications. Go to the following website and click **My notifications** for more information about how to register for support notifications:

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

You can also find product updates and support notifications if you use the IBM Support Portal website at:

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

Chapter 1. Preparing for installation

The following information helps you to prepare for the successful installation of the Storage Manager software.

- “Storage Manager software”
- “Supported controller firmware” on page 2
- “Types of installation configurations” on page 2
- “Setting up controller addresses for software installation” on page 7

Introduction

Lists the operating systems on which the Storage Manager is supported for DS3000, DS4000, and DS5000 Storage Subsystems, and also the operating systems that are supported when the storage subsystems are attached to it.

The IBM System Storage DS Storage Manager consists of a set of client and host tools that you can use to manage the IBM DS3000, DS4000, and DS5000 Storage Subsystems from a management station.

The Storage Manager is supported on the following operating systems:

- AIX
- Windows Server 2003, Windows Server 2008, and Windows Server 2012
- Linux (RHEL and SLES)
- HP-UX
- Solaris

The DS3000, DS4000, and DS5000 Storage Subsystems are also supported when they are attached to Apple Mac OS, VMware ESX Server, and System p Virtual IO Server (VIOS) hosts, as well as on i5/OS as a guest client on VIOS. IBM does not provide host software for these operating systems. You must install IBM DS Storage Manager on a management station which has one of the operating systems listed above installed.

Information about i5/OS support can be found at the following website:

www.ibm.com/systems/i/os/

For additional information, see the System Storage Interoperation Center at the following website:

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

Storage Manager software

Lists the tasks you can perform using the Storage Manager software, including accessing premium features.

Storage Manager is used to configure, manage, and troubleshoot storage subsystems. It is used primarily to configure disk pools or RAID arrays and logical drives, assign logical drives to hosts, replace and rebuild failed disk drives, expand the size of the disk pools, arrays and logical drives, and convert from one RAID

level to another. Storage Manager enables troubleshooting and management tasks, such as checking the status of the storage subsystem components, updating the firmware of the RAID controllers, and managing the storage subsystem. Finally, the Storage Manager offers access to premium features such as FlashCopy®, VolumeCopy, and Enhanced Remote Mirroring.

For the latest firmware versions that are supported by each storage subsystem model, see the readme file for your operating system.

Storage Manager software components

Lists the components of the Storage Manager software and the differences depending on the operating system.

Storage Manager includes the following client software components.

Note: Storage Manager components might vary depending on the operating system. For Storage Manager version 10.77.xx.xx and later, the Microsoft MPIO DSM installer is separate from the Storage Manager installer that lets you install the components listed below. However, the Storage Manager installer and the Microsoft MPIO DSM installer are bundled in a single code package. This code package is available on the IBM support portal.

SMruntime software

Storage Manager Java™ compiler

SMesm software

Storage Manager ESM firmware delivery package

SMclient software

Storage Manager client package

SMagent software

Storage Manager agent package

SMutil software

Storage Manager utility package

Supported controller firmware

All controller firmware versions are available at no cost on the IBM website.

To achieve the highest level of compatibility and error-free operation, make sure that the controller firmware for your storage subsystem is the latest firmware version for the storage subsystem model.

Important: If the Storage Manager version is 10.84.xx.xx or later, controller firmware must be 6.50.xx.xx or later.

For detailed information about how to download the most current firmware version level, see “Downloading controller firmware, NVSRAM, ESM firmware” on page 43.

Types of installation configurations

Defines a Network configuration and a Direct-attached or SAN-attached configuration.

A management station can be either of the following configurations:

Network configuration (out-of-band)

A remote system, connected to an Ethernet network, that is used to manage one or more storage subsystems.

Direct-attached or SAN-attached configuration (in-band or out-of-band)

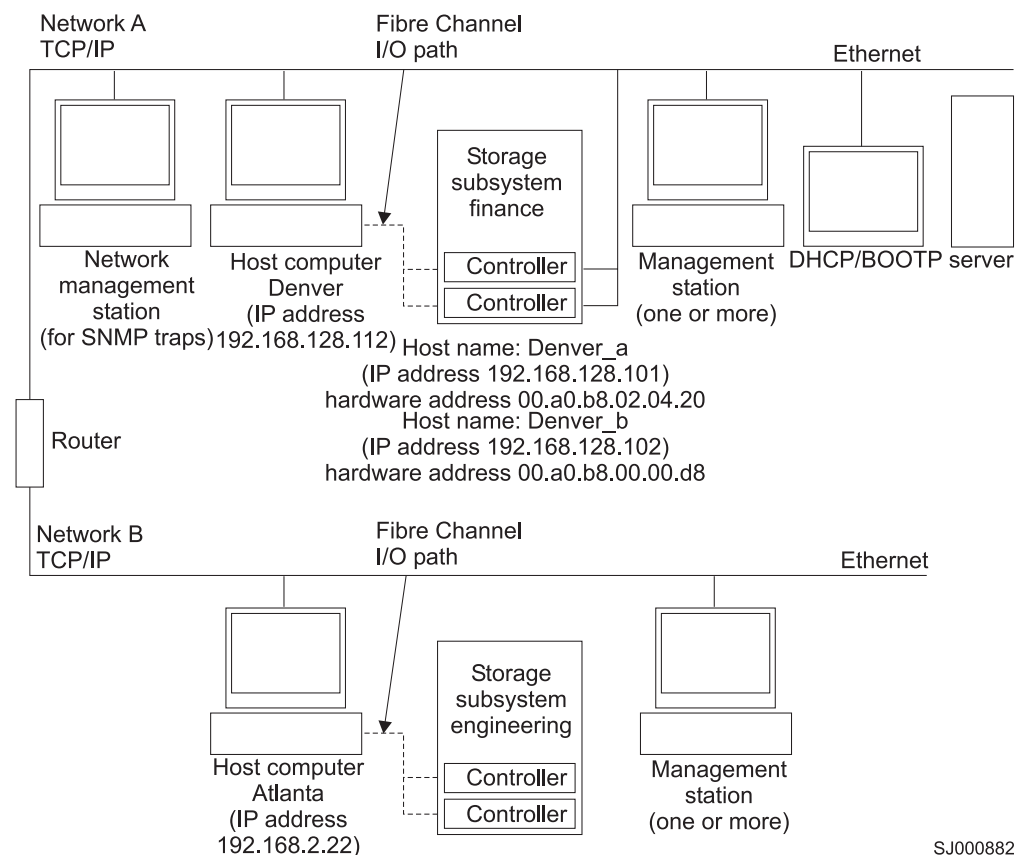
A host that is connected to a storage subsystem through a Fibre Channel, iSCSI, or SAS input/output (I/O) path. This host can use either the I/O path (in-band) or its Ethernet network ports (out-of-band).

Network configuration

Informs the network-related tasks you must perform before installing the Storage Manager.

Before you begin installing the Storage Manager software, ensure that the network components are set up and operating properly and that you have all of the host and controller information that is necessary for the correct operation of the software.

Note: When you connect the storage subsystem to an Ethernet switch, set the switch port settings to auto-negotiate.



SJ000882

Figure 1. Sample network using network-managed and host-agent-managed storage subsystems

Reviewing a sample network configuration

Reviews the components of a network managed storage subsystem and a host-agent-managed storage subsystem.

Figure 1 on page 3 shows an example of a network that contains both a network managed storage subsystem (Network A) and a host-agent-managed storage subsystem (Network B).

Network-managed storage subsystem: Network A is a network-managed storage subsystem. Both the management station and the storage subsystem are connected to the Ethernet network. Network A contains the following components:

- A DHCP/BOOTP server
- A network-management station for Simple Network Management Protocol (SNMP) traps
- A host that is connected to a storage subsystem through a Fibre Channel I/O path
- A management station that is connected by an Ethernet cable to the storage subsystem controllers

Note: If the controller static TCP/IP addresses or default TCP/IP addresses are used, you do not have to set up the DHCP/BOOTP server.

Host-agent-managed storage subsystem: Network B is a host-agent-managed storage subsystem. You can manage the storage subsystem using the same path, fibre channel or SAS, that a host server uses to send I/O between the host and the storage subsystem using the Storage Manager agent software that is installed in the host server. The storage subsystem Ethernet management ports are not normally connected to the Ethernet network.

The Storage Manager agent requires a special LUN be assigned/mapped to the host partition. This LUN is referred to as Access or UTM LUN and is used by the Storage Manager agent and the controller to pass management information. This LUN is automatically assigned to the host partition as LUN 31 the first time a logical drive is assigned to a host partition. Because the LUN has reduced the maximum number of LUN/logical drives that can be assigned to a host by one, use the storage subsystem management GUI to unassign it if the storage subsystem is managed solely via out-of-band method.

Note: The storage subsystem can be managed in-band and out-of-band at the same time.

Network B contains the following components:

- A host that is connected to a storage subsystem through a supported I/O path
- A management station that is connected by an Ethernet cable to the host computer

Setting up a management station

Defines a management station and informs about the tasks it performs.

The *management station* is the server that is responsible for managing all of, or a portion of, a storage network. It communicates with the network management agents in the managed nodes, using a network management protocol such as Simple Network Management Protocol (SNMP).

Storage management commands are sent to the storage subsystem controllers, where the controller firmware validates and runs the commands and then returns status and configuration information to the client software.

Setting up a network-managed (out-of-band) configuration

Informs you how to set up a network-managed (out-of-band) configuration.

The following steps provide an overview of the tasks that are required to set up the network for installation of a network-managed (out-of-band) configuration:

Important: A maximum of eight management stations can concurrently monitor an out-of-band-managed storage subsystem. This limit does not apply to servers that manage the storage subsystem through the in-band-management method.

1. Install all of the hardware components (host computers, storage subsystems, and cables) that you want to connect to the network. For more information about installing hardware components, see the documentation that came with the hardware components.
2. Establish a naming convention for the storage subsystems that will be connected to the network.
3. Record the storage subsystem names and management types.

Note: Throughout the remaining steps, you must record some information for future use, such as the hardware Ethernet and IP addresses.

4. Determine the hardware Ethernet MAC address for each controller in storage subsystems connected to the network. If you are using a default controller IP address, go to step 6. Otherwise, obtain the TCP/IP address and host name for each of the controllers in the storage subsystems on the network from the network administrator.
5. Set up the DHCP/BOOTP server to provide network configuration information for a specific controller. If you are using static controller IP addresses, skip this step.
6. Verify that the TCP/IP software is installed.
7. Set up the host or domain name server (DNS) table.
8. Turn on the power to the devices that are connected to the network.

Setting up a host-agent-managed (in-band) configuration

Informs you how to set up a host-agent-managed (in-band) configuration.

The following steps provide an overview of the tasks that are required to set up an installation of a host-agent-managed (in-band) configuration:

1. Install all of the hardware components (host computers, storage subsystems, and cables) that you want to manage. For more information about installing hardware components, see the documentation that came with the hardware components. The host computer must have configured I/O connections to the storage subsystem (for example, the host must have an operating system installed with the applicable device driver for the host bus adapters).
2. Install the Storage Manager host software and the Storage Manager agent software.
3. Establish a naming convention for the storage subsystems that will be connected to the network.
4. Record the storage subsystem names and management types.

Note: Throughout the remaining steps, you must record some information for future use, such as the hardware Ethernet and IP addresses.

5. Obtain the IP address and host name of the host computer on which the host-agent software will run from the network administrator.

Note: SAgent is part of the Storage Manager software package and is required on the host that is connected to the storage subsystem through any of the supported interfaces.

6. Verify that the TCP/IP software is installed.
7. Turn on the power to the devices that are connected to the network.

Note: Even though you do not connect the host and the storage subsystems management Ethernet ports to the network, the host still uses TCP/IP to communicate with the host-agent. The host-agent communicates with the controller over the Fibre Channel connection through the *access volume*.

Direct-attached and SAN-attached configurations

Storage Manager supports in-band management of storage subsystems in direct-attached configurations or in a SAN environment through switches.

Setting up a direct-attached configuration

Informs you how to directly connect a storage subsystem to the Storage Manager.

Important: Storage subsystems with iSCSI ports do not support direct-attached connections from the host systems to the storage subsystem iSCSI ports.

Before you begin, verify that:

- You can connect one or two servers to the storage subsystems.
- No external switches or external Fibre Channel hubs are being used.
- See the *Installation and User's Guide* for your storage subsystem for more information.

Complete the following steps to set up a direct-attached configuration:

1. Connect the HBAs to each controller port of the storage subsystem.
2. Use the Storage Manager automatic discovery feature to make sure that the storage subsystem is discovered.

Setting up a SAN-attached configuration

Informs you how to connect the storage subsystems to the Storage Manager in a SAN.

A SAN-attached configuration can consist of Fibre Channel, SAS, or iSCSI connections.

If you use Fibre Channel HBAs in your SAN-attached configuration, the HBA and the storage subsystem host port connections should be isolated in fabric zones to minimize the possible interactions between the ports in a SAN fabric environment. Multiple storage subsystems can be configured to the same set of HBAs through a Fibre Channel, SAS, or Ethernet switch. For more information about Fibre Channel zoning schemes, see “Connecting HBAs in a Fibre Channel switch environment” on page 119. Similar zoning schemes can be implemented with SAS and Ethernet switches also.

Attention: A single-HBA configuration can result in loss of data access in the event of a path failure. If you have a single HBA in a SAN-attached configuration, both controllers in the storage subsystem must be connected to the HBA through a switch, and both controllers must be within the same SAN zone as the HBA.

Complete the following steps to set up a SAN-attached configuration:

1. Connect the HBAs to the switch or switches.
2. Connect the storage subsystems to the switch or switches.

3. Set the required zoning or VLANs on the Fibre Channel switches or Ethernet switches, if applicable.
4. Use the Storage Manager automatic discovery feature to make sure that the storage subsystem is discovered.

Setting up controller addresses for software installation

How you plan to manage the storage subsystems determines where you must install the software components. Before you can install software components, you must assign IP addresses for the storage controllers.

Note:

1. The controllers must be connected to a LAN port that is set to auto-negotiate the data rate. The controllers do not function properly when they are connected to a switch port that is set for a fixed rate.
2. To manage storage subsystems through a firewall, configure the firewall to open port 2463 to TCP data.

Setting up IP addresses for storage subsystem controllers

Lists components of a DHCP or BOOTP server and network and explains how IP addresses are assigned to storage subsystem controllers.

Complete the following procedures after you install SMruntime and SMclient, as described in the installation section for your host operating system.

You must set up a DHCP or BOOTP server and network with the following components:

- A DHCP or BOOTP server
- A network-management station for Simple Network Management Protocol (SNMP) traps
- A host that is connected to a storage subsystem through a Fibre Channel I/O path
- A management station that is connected by an Ethernet cable to the storage subsystem controllers

Note: You can avoid DHCP/BOOTP server and network tasks by assigning static IP addresses to the controller. If you do not want to assign static TCP/IP addresses with the Storage Manager, using the storage subsystem default TCP/IP addresses as shown in “Assigning static TCP/IP addresses to a storage subsystem using factory-default management port TCP/IP address” on page 9, establish an in-band management connection to the storage subsystem and change the management port IP address in Subsystem Management window.

If a controller has two management ports, the same gateway address is shared between the two ports. The most recently obtained or supplied gateway address is used for both ports. Therefore, it is possible to lose access on one port as a result of changing the configuration on the other port. If both ports are manually configured, the most recently supplied gateway address will be used. If one port is manually configured and DHCP is enabled on the other port, the most recently supplied or obtained gateway address will be used. Generally, this is the gateway address supplied by the DHCP server unless the manual configuration for the other port is changed. In this case, the gateway address should be set to the value provided by the controller, which should match the gateway address obtained from the DHCP server. If DHCP is enabled on both ports, the DHCP servers

attached to the two ports should be configured to supply the same gateway address. If the DHCP servers apply different gateway addresses, the most recently obtained gateway address will be used for both ports.

Any changes to remote login access affect both ports. In other words, if remote login access is enabled or disabled on one port, it is also enabled or disabled on the other port. As with the gateway address, the most recent configuration applied for remote login applies to both ports. For example, if remote login access is manually enabled on port 1, it will also be enabled for port 2. If a DHCP server subsequently supplies configuration parameters for port 2 that includes disabling remote login access, it will be disabled for both ports.

If a controller has two management ports, the two Ethernet ports must be on different subnets. If both ports are on the same subnet, or if they have the same network address (the logical AND of the IP address and the subnet mask), Subnet Configuration Error event notification will occur.

Setting up an IP address with the DHCP/BOOTP server

This topic describes the steps to set up the DHCP/BOOTP server and network.

Complete the following steps to set up the DHCP/BOOTP server and network:

1. Get the MAC address from each storage subsystem controller. (See the “Identifying Ethernet MAC addresses” procedure.)
2. Complete whichever of the following steps is applicable for your server:
 - On a DHCP server, create a DHCP record for each of the MAC addresses. Set the lease duration to the longest time possible.
 - On a BOOTP server, edit the bootptab file to add the entries that associate the MAC address tab with the TCP/IP address.
3. Connect the DS3000, DS4000, or DS5000 storage subsystem Ethernet ports to the network.
4. Boot the storage subsystem.

Identifying Ethernet MAC addresses

Explains the hardware Ethernet medium access control (MAC) addresses that the controllers have, with the format and an example.

To manage your storage subsystem with the direct-management method, you must identify the hardware Ethernet medium access control (MAC) address for each controller.

Every storage subsystem has a label with the hardware Ethernet MAC address number. The number has the format *xx.xx.xx.xx.xx.xx*, where *x* represents a letter or a number. For example, an Ethernet MAC address might be 00.a0.b8.20.00.d8.

Instructions and label locations for particular storage subsystems are listed in the following sections.

Identifying the Ethernet MAC addresses on a DS4800, DS5100, or DS5300

storage subsystem: The machine type, model number, and serial number are on top of each RAID controller unit. The MAC addresses are near the Ethernet ports on each RAID controller.

Note: You can access the controllers from the back of a DS4800, DS5100, or DS5300 chassis.

Identifying the Ethernet MAC addresses on a DS3000, DS3500, DCS3700, DS3950, DS4200, DS4700, DS5020 storage subsystem, or a DCS3700 storage subsystem with Performance Module Controllers: The MAC addresses on these storage subsystems are near the Ethernet ports on each RAID controller.

Note: You can access the controllers from the back of the storage subsystem chassis.

Identifying the Ethernet MAC addresses on a DS4400, or DS4500 storage subsystem: To identify the hardware Ethernet MAC addresses for DS4400 and DS4500 storage subsystems, complete the following steps:

1. Remove the front bezel from the storage subsystem and carefully pull the bottom of the bezel out to release the pins. Then slide the bezel down.
2. On the front of each controller, look for a label with the hardware Ethernet MAC address. The number is in the form *xx.xx.xx.xx.xx.xx* (for example, 00.a0.b8.20.00.d8).
3. Record each Ethernet MAC address.
4. To replace the bezel, slide the top edge under the lip on the chassis. Then push the bezel bottom until the pins snap into the mounting holes.

Identifying the Ethernet MAC addresses on a DS4100, or DS4300 storage subsystem: To identify the hardware Ethernet MAC address for machine types 3542 (DS4100, or 1722 (DS4300), complete the following steps:

1. Locate the Ethernet MAC address at the back of the unit, under the controller Fibre Channel host ports. The number is in the form *xx.xx.xx.xx.xx.xx* (for example, 00.a0.b8.20.00.d8).
2. Record each Ethernet MAC address.

Assigning static TCP/IP addresses to a storage subsystem using factory-default management port TCP/IP address

Explains how to assign static TCP/IP addresses to the storage subsystem controllers using the factory defaults.

Complete the following steps to assign static TCP/IP addresses to the storage subsystem controllers, using the default TCP/IP addresses that are assigned to the controllers when they are manufactured:

1. Make a direct-management connection to the storage subsystem, using the default TCP/IP addresses for the controllers. To find the default TCP/IP addresses for your storage subsystem, see the *Installation and User's Guide* that was shipped with the hardware installed on controller management port labeled #1.
 - Controller A: 192.168.128.101
 - Controller B: 192.168.128.102
 - Subnet Mask: 255.255.255.0

Note: For a storage subsystem that has two Ethernet ports per controller (such as the DCS3700, DCS3700 with Performance Module Controllers, DS5020, DS3500, DS5100 and DS5300), use the Ethernet port that is labeled #2. The default IP addresses of the second Ethernet port are as follows:

- Controller A: 192.168.129.101
- Controller B: 192.168.129.102
- Subnet Mask: 255.255.255.0

2. Start SMclient. The Enterprise Management window opens.
3. In the Enterprise Management window, click the name of the default storage subsystem. The Subsystem Management window opens.
4. In the Subsystem Management window, right-click the controller icon and select **Change > Network Configuration** in the menu. The Change Network Configuration window opens.
5. In the Change Network Configuration window, click the **Controller A** and **Controller B** tabs and type the new TCP/IP addresses in the applicable fields. Click **OK**.
6. Close the Subsystem Management window, wait 5 minutes, and delete the default storage subsystem entry in the Enterprise Management window.
7. Add a new storage subsystem entry in the Enterprise Management window, using the new TCP/IP address.

Assigning static TCP/IP addresses storage subsystem using an in-band management connection

For a host that is connected to a storage subsystem through a Fibre Channel I/O path, complete these steps to assign static TCP/IP addresses to the storage subsystem controllers by way of the host that has Fibre Channel connectivity to the storage subsystem Fibre Channel host port.

To complete this procedure, you must have the following components:

- A host that is connected to a storage subsystem through a Fibre Channel I/O path
- A management station that is connected by an Ethernet cable to the storage subsystem controllers

Note: You cannot perform in-band management with a host that has iSCSI connections to the storage subsystem until it is configured. Instead, use other methods in this section to assign static TCP/IP address to the storage subsystem controller management ports.

1. Install the DS Storage Manager client software in the host, and make sure that the SMagent software is installed and running.
2. Start the DS Storage Manager client software. The Enterprise Management window opens.
3. Add the storage subsystem to the Enterprise Management domain using the IP address of the host that has Fibre Channel connectivity to the storage subsystem.
4. In the Enterprise Management window, click on the name of the newly-discovered storage subsystem. The Subsystem Management window opens.
5. In the Subsystem Management window, right-click the Controller icon and select **Change > Network Configuration** in the drop-down menu. The Change Network Configuration window opens.
6. In the Change Network Configuration window, click on the Controller A and Controller B tabs and type the new TCP/IP addresses in their applicable fields.
7. Click **OK**.
8. Close the Subsystem Management window.
9. Wait at least five minutes.

10. Delete the existing storage subsystem entry in the Enterprise Management window.
11. If applicable, change the IP address of the Ethernet port in the management station to a value that is on the same TCP/IP subnet as the controller Ethernet port IP addresses that you just assigned.
12. Exit DS Storage Manager.
13. Restart.
14. Make Ethernet cabling to the controller management ports.
15. Add a new storage subsystem entry in the Enterprise Management window, using the new assigned IP addresses.

Assigning static TCP/IP addresses using the storage subsystem controller serial port Service Interface

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

To complete this procedure, you must have the following components:

- A null modem cable with DB-9 female connectors on both ends, used to connect the host serial port to the controller serial port.
- A terminal emulation software application, such as Procomm or Microsoft Windows Hyperterm, installed in the host system.

Note:

1. The terminal session setting must have these values: 38400 BAUD; 8 data bits; 1 stop bit; no parity.
2. If the controller BAUD rate setting is different from the terminal setting, send a "break" character to cause the controller to switch to the next available BAUD rate setting. Repeat sending the "break" character until the "Press space to set the BAUD rate" message is displayed.
- Controller firmware version 7.77.xx.xx or higher and its associated NVSRAM files installed.

Complete the following steps to view and assign new IP address to the controller management port:

1. Press **Enter**. If this character (->) is displayed, type **Exit** and press **Enter**. Otherwise, continue to next step.
2. In the terminal emulator session, send the "break" character. For example, use **CNTL+BREAK** for Microsoft Windows Hyperterm or **ALT+B** for Procomm.
3. Enter the uppercase letter **S** and press **Enter** when the following message is displayed: Press within 5 seconds: for <S> Service Interface, <BREAK> for baud rate.
4. Enter the password **DSStorage** (case sensitive) within 60 seconds of when this message is displayed: Enter the password to access the Service Interface (60 second timeout).

Note: If the controller does not have controller firmware version 7.77.xx.xx or higher and its associated NVSRAM files installed, this password will not be accepted, and you must follow one of the two methods to change the IP configuration of controller Ethernet ports. See "Assigning static TCP/IP addresses to a storage subsystem using factory-default management port

TCP/IP address” on page 9 and “Assigning static TCP/IP addresses storage subsystem using an in-band management connection” on page 10 for more information.

5. Enter 1 or 2 to display or change the IP configuration when the following menu is displayed:

Service Interface Main Menu

=====

- 1) Display IP Configuration
- 2) Change IP Configuration
- 3) Reset Storage Array Administrator Password
- Q) Quit Menu

If option 2 is chosen, follow the prompt to set the IP configuration for the port that you selected. You must reboot the controller for the settings to take effect.

Note: You must perform these steps on both controllers.

Chapter 2. The Storage Manager interface

Describes the basic layout of the Storage Manager software, where Storage Manager version is 10.84.xx.xx and controller firmware version is 7.84.xx.xx.

This chapter describes the basic layout of the Storage Manager software on a subsystem, where Storage Manager version is 10.84.xx.xx and controller firmware version is 7.84.xx.xx.

Storage Manager has two windows that provide management functionality and a graphical representation of your storage subsystems: the Enterprise Management window and the Subsystem Management window.

Use the Enterprise Management window to add the storage subsystems that you want to manage and monitor. Through the Enterprise Management window, you receive alert notifications of critical errors that are affecting the storage subsystems. If you are notified in the Enterprise Management window that a storage subsystem has a non-optimal status, you can open the Subsystem Management window for the affected storage subsystem to see detailed information about the storage subsystem condition.

Important: Depending on the version of your Storage Manager and controller firmware, the views, menu options, and functionality might differ from the information that is presented in this document. For information about available functionality, see online help topics of the Storage Manager.

Enterprise Management window

Lists the tasks you can perform using the Enterprise Management window.

The Enterprise Management window is the first window that opens when you start Storage Manager. Use the Enterprise Management window to complete the following management tasks:

- Discover in-band hosts and out-of-band storage subsystems automatically on your local subnetwork
- Manually add and remove in-band hosts and storage subsystems
- Monitor the health of the storage subsystems and report a high-level status with the applicable icon
- Configure alert notifications through email or Simple Network Management Protocol (SNMP) and report critical events to the configured alert destinations

Note: A local configuration file stores all of the information about storage subsystems that you have added and any email destinations or SNMP traps that you have configured.

- Open the applicable Subsystem Management window for a selected storage subsystem to perform detailed configuration and management operations
- Execute scripts to perform batch management tasks on a particular storage subsystem

Note: For example, you might run scripts to create new logical drives or to download new controller firmware. For more information about executing

scripts, see the *IBM System Storage DS Storage Manager Command Line Interface and Script Commands Programming Guide* or the online help topics in the Enterprise Management window.

- Schedule save or automatically save a copy of the support data when the client monitor process detects a critical event.
- Upgrade controller firmware between major controller firmware versions (for example, upgrading controller firmware from version 6.xx.xx.xx to 7.xx.xx.xx). This functionality is the same as that of the stand-alone IBM System Storage Controller Firmware upgrade tool. This tool is integrated into the IBM DS Storage Manager client version 10.50.xx.xx and later.
- Schedule collection of all support information on one or all of the storage subsystems at definite intervals.
- Retrieve firmware inventory of all storage subsystems.

To display or hide the Tool and Status bars, select View from the menu and select or clear the Tool or Status options.

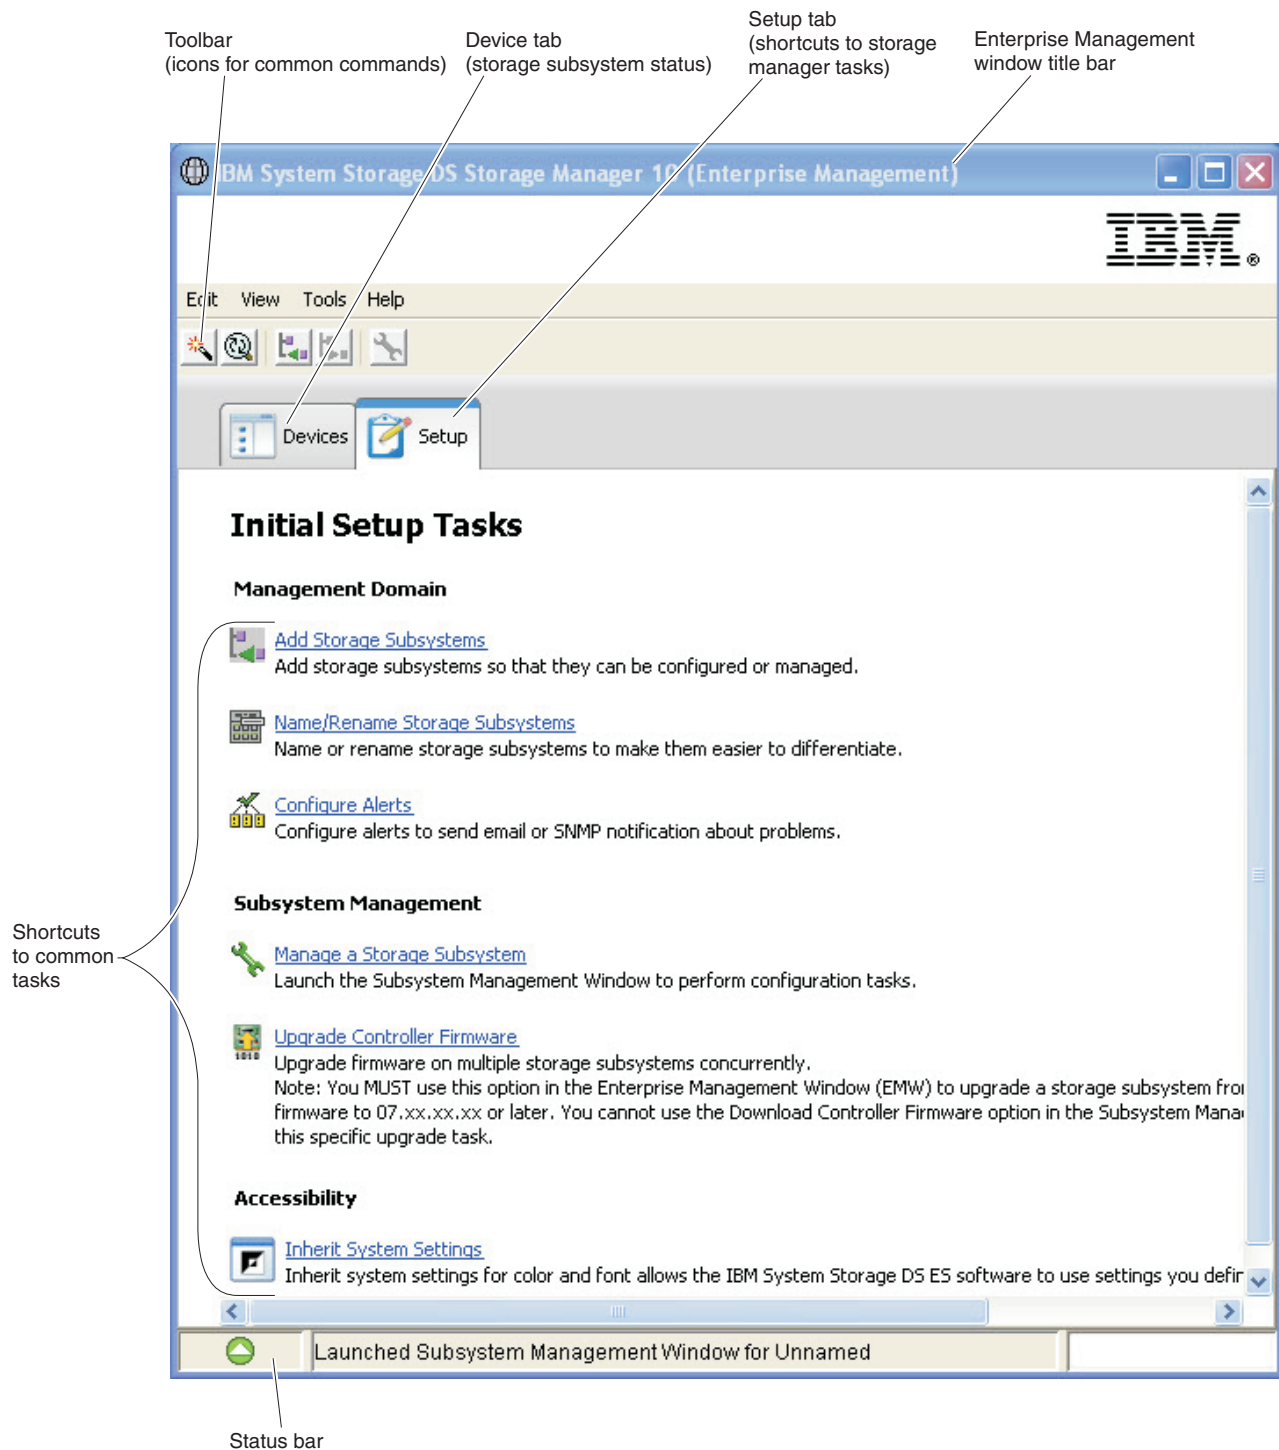


Figure 2. Parts of the Enterprise Management window

Using the Devices tab

Describes the two views - tree view and table view that the Devices tab in the Enterprise Management window and the tasks you can perform in each view.

The **Devices** tab in the Enterprise Management window presents two views of the storage subsystems that are managed by the management station: a tree view and a table view.

Tree view

The tree view provides a hierarchical view of the nodes in the storage subsystem. The tree view shows two types of nodes:

- Discovered Storage Subsystems
- Unidentified Storage Subsystems

The Discovered Storage Subsystems node and the Unidentified Storage Subsystems node are child nodes of the Management Station node.

The Discovered Storage Subsystems node has child nodes that represent the storage subsystems that are currently managed by the management station. Each storage subsystem is labeled with its machine name and is always present in the tree view. When storage subsystems and hosts with attached storage subsystems are added to the Enterprise Management window, the storage subsystems become child nodes of the Discovered Storage Subsystems node.

Note: If you move the mouse over the Discovered Storage Subsystems node, a tooltip appears displaying the controller IP address.

The Unidentified Storage Subsystems node shows storage subsystems that the management station cannot access because of network connection problems, turned off subsystem, or a non-existent name.

You can perform these actions on the nodes in the tree view:

- Double-click the Management Station node and the Discovered Storage Subsystems node to expand or collapse the view of the child nodes.
- Double-click a storage subsystem node to launch the Subsystem Management window for that storage subsystem.
- Right-click the Discovered Storage Subsystems node to open a menu that contains the applicable actions for that node.

The right-click menu for the Discovered Storage Subsystems node contains these options:

- **Add Storage Subsystem**
- **Automatic Discovery**
- **Refresh**
- **Collect Support Data**
 - automatically
 - create/edit schedule

The right-click menu for the storage subsystem nodes contains these options:

- Manage storage subsystem
- Locate storage subsystem
- Execute script
- Load storage subsystem configuration
- Upgrade controller firmware
- Refresh

- Remove storage subsystem
- Remove management connection
- Configure Alerts
- Collect support data automatically
- Create/Edit Collect support data schedule
- Rename the storage subsystem
- Comment

These options are also included with the other options in the **Edit** and **Tools** menu options. For more information, see the *Using the Enterprise Management window* online help topic.

Table view

In the table view, each storage subsystem is a single row in the table. The columns in the table view show data about the managed storage subsystem.

Table 1. Data shown in the table view

Column	Description
Name	The name of the managed storage subsystem Note: If the managed storage subsystem is unnamed, the default name is Unnamed.
Type	The type of managed storage subsystem, represented by an icon
Status	An icon and a text label that report the true status of the managed storage subsystem
Management Connections	The following connection types are possible: <ul style="list-style-type: none"> • Out-of-Band: this storage subsystem is an out-of-band storage subsystem. • In-Band: this storage subsystem is an in-band storage subsystem that is managed through a single host. • Out-of-Band, In-Band: this storage subsystem is a storage subsystem that is both out-of-band and in-band. Click Details to see more information about any of these connections.
Comment	Any comments that you have entered about the specific managed storage subsystem

Sort the rows in the table view in ascending order or descending order by either clicking a column heading or by selecting one of these menu options:

- **View > By Name**
- **View > By Status**
- **View > By Management Connection**
- **View > By Comment**

To change the way that managed storage subsystems appear in the table view, complete one of the following actions:

- To show all of the known managed storage subsystems in the table view, select the Management Station node.

- To show only that storage subsystem in the table view, select a storage subsystem node in the tree view.

Note: Selecting an Unidentified node in the tree view shows an empty table view.

Showing Managed Subsystems in the Table View

You can change the way managed storage subsystems appear in the Table view.

- Select the storage manager node to show all of the known managed storage subsystems in the Table view.
- Select a Discovered Storage Subsystem node or an Undiscovered Storage Subsystem node in the Tree view to show any storage subsystem that are attached to that specific host in the Table view.

Note: If you have not added any storage subsystems, the Table view is empty.

- Select a storage subsystem node in the Tree view to show only that storage subsystem in the Table view.

Note: Selecting an Unidentified node in the Tree view shows an empty Table view.

Adding and Removing a Storage Subsystem

Do one of the following in the Storage Manager to add a storage subsystem.

Table 2. Adding a Storage Subsystem

Location	Procedure
Tree view	Right-click the root node from the Tree view, select Add Storage Subsystem from the pop-up menu
Toolbar	Click the icon to add the Storage Subsystem
Edit menu	Select Edit>Add Storage Subsystem
Setup tab	Select Add Storage Subsystem

To remove a storage subsystem, do one of the following in the Storage Manager. Removing a subsystem removes the icon only, and does not delete the subsystem. You can select more than one subsystems at a time.

Table 3. Removing Storage Subsystems

Location	Procedure
Tree View	Right-click the storage subsystem that you want to remove from the Tree view, and select Remove > Storage Subsystem from the pop-up menu
Toolbar	Select the storage subsystem that you want to remove from the Tree view or Table view, and click the icon to remove the storage subsystem
Edit menu	Select the storage subsystem that you want to remove from the Tree view or Table view, and select Edit > Remove > Storage subsystem

Removing Multiple Storage Subsystems simultaneously

If you are managing numerous storage subsystems you can remove two or more continuous or non-continuous storage subsystems at the same time by using the Table view.

Note: You remove only the icon from the Tree or Table view and you do not remove the storage subsystems.

Table 4. Simultaneously removing multiple subsystems

Location	Procedure
Tree View	Click the Discovered Storage Subsystem node to display the storage subsystems you are managing
Table View	To remove continuous storage subsystems, click the first storage subsystem you want to remove. Hold the Shift key and click the last storage subsystem you want to remove. With the selected storage subsystems highlighted in the Table view, right-click and select Remove from the pop-up menu. To remove non-continuous storage subsystems, hold the Control key and click the storage subsystems you want to remove. With the selected storage subsystems highlighted in the Table view, right-click and select Remove from the pop-up menu.

Using the Setup tab

Describes the Setup tab in the Enterprise Management window and the tasks you can perform.

The Enterprise Management window **Setup** tab is a gateway to tasks that you can perform when you set up a storage subsystem. Use the Enterprise Management window Setup tab to perform the following tasks:

- Add a storage subsystem
- Name or rename a storage subsystem
- Configure an alert
- Open the Subsystem Management window to manage a storage subsystem
- Upgrade controller firmware
- Open the Inherit Systems Settings window

Subsystem Management window

Describes how the Subsystem Management window is launched and the tasks you can perform using this window.

The Subsystem Management window is Java technology-based software that is launched from the Enterprise Management window. Each Subsystem Management window provides management functions for a single storage subsystem. You can have more than one Subsystem Management window open to manage different storage subsystems. The Subsystem Management window includes the following functions:

- Access storage subsystem options, such as locating a storage subsystem, configuring a storage subsystem, renaming a storage subsystem, or changing a password
- Configure disk pools or arrays and thin standard or thin logical drives from your storage subsystem capacity, define hosts and host groups, and grant host or host group access to sets of standard or thin logical drives called storage partitions
- Monitors the health of storage subsystem components and reports a detailed status using applicable icons
- Access the applicable recovery procedures for a failed logical component or a failed hardware component
- View the event log for the storage subsystem
- View profile information about hardware components, such as controllers and drives and get a physical view of the drives in the hardware enclosures
- Access controller-management options, such as changing ownership of logical drives or placing a controller online or offline
- Access drive-management options, such as assigning hot spares and locating the drive
- Monitor storage subsystem performance
- Configure copy services like Enhanced Flashcopy, Flashcopy, VolumeCopy, and Remote Mirroring

If the storage subsystem has controller firmware version 7.70.xx.xx, its Subsystem Management window cannot be opened unless a strong password is provided. A strong password must be between 8 and 30 characters and contain at least one number, one lower-case letter, one upper-case letter, and one non-alphanumeric character (for example, < > ! @ + #). Spaces are not permitted, and it is case-sensitive.

In storage subsystems with controller firmware other than 7.70.xx.xx, you are prompted to provide this password, if none is specified for the storage subsystem, whenever you attempt to open a Subsystem Management window for this storage subsystem. IBM recommends creating a subsystem management password to prevent unauthorized changes to the Subsystem Management configuration.

Opening the Subsystem Management window

Lists the four methods to open the Subsystem Management window

To open a Subsystem Management window from the Enterprise Management window, perform one of the following actions:

- Click the **Devices** tab, and double-click the name of the storage subsystem that you want to manage.
- Click the **Devices** tab, right-click the name of the storage subsystem that you want to manage, and select **Manage Storage Subsystem**.
- Click the **Devices** tab, and select **Tools > Manage Storage Subsystem**.
- Click the **Setup** tab, and select **Manage Storage Subsystem**. In the Select Storage Subsystem window, select the name of the storage subsystem that you want to manage, and click **OK**.

You can manage only a single storage subsystem within a Subsystem Management window. However, you can open more than one Subsystem Management window from the Enterprise Management window to simultaneously manage multiple storage subsystems.

The Subsystem Management window provides the following options for managing your storage subsystem.

The screen layout and the menu options of the Subsystem Management window have changed from IBM System Storage DS Storage Manager version 10.83 and later, in conjunction with controller firmware version 7.83.xx.xx and later. All of the subsystem management functions that were implemented in the earlier versions are still valid. However, the menu options might be different. Explore the software and refer to the online help to acquaint yourself with it.

Using the Summary tab

Lists information you can view on the summary tab of the Subsystem Management window.

The **Summary** tab in the Subsystem Management window shows information about the storage subsystem. The Summary tab also includes links to the Storage Subsystem Profile window, relevant online help topics, and the storage concepts tutorial. The link to the Recovery Guru window is also shown, when the storage subsystem needs attention.

On the Summary tab, you can view this information:

- Status of the storage subsystem
- Version information for IBM DS Storage Manager software and firmware controller
- Capacity of the storage subsystem
- Disk pools and arrays, Logical drives, and Copy Service configurations like FlashCopy, VolumeCopy and enhanced remote mirroring in the storage subsystem
- Hosts, the mappings, and the storage partitions in the storage subsystem
- Number of premium features available, active, enabled, or disabled for your storage subsystem
- Hardware components in the storage subsystem
- Online documentation available for learning about your storage subsystem

Using the Storage and Copy Services tab

Informs about the tree-structured view of the logical nodes that the Storage & Copy Services tab provides.

The **Storage & Copy Services** tab provides a tree-structured view of the logical nodes. Click the plus (+) sign or the minus (-) sign adjacent to a node to expand or collapse the view. Right-click a node to open a menu that contains the applicable actions for that node.

Nodes on the Logical tab

The storage subsystem, or root node, has the types of child nodes that are shown in the following table.

Table 5. Nodes on the Logical tab

Child nodes of the root node	Description of the child nodes
All Logical Objects	This node lets you view information about all the logical objects that comprise your storage subsystem. Use the Object Type drop-down menu in the View pane to select a particular object type. This is a useful way to view the status and capacity information for a disk pool or an array, or view all the repository logical drives that are associated or not associated with a base logical drive used with the Flashcopy Image, Enhanced Flashcopy Image, and Consistency Group premium features.
Total Unconfigured Capacity	This node represents the sum of the capacity of all unassigned drives that are not in a disk pool or an array.
Unconfigured Capacity	This node represents the storage subsystem capacity that is not configured into an array. Multiple Unconfigured nodes appear if your storage subsystem contains drives with different media types (hard disk drives or solid state drives) and different interface types. Each drive type has an associated Unconfigured Capacity node shown under the Total Unconfigured Capacity node if unassigned drives are available in a drive enclosure.
Disk Pools	<p>The IBM Storage Manager displays a Disk Pools node if one or more disk pools have been configured for your storage subsystem. Expand the Disk Pool node to see the individual disk pools. If the Flashcopy premium feature is enabled, you can have the flashcopy image child nodes. The disk pool node has several types of child nodes:</p> <ul style="list-style-type: none"> • Logical Drive — This node represents a configured and defined logical drive (either of standard or thin logical drive). Multiple Logical drive nodes can exist under a Disk Pools node. • Free Capacity — This node represents a region of capacity that you can use to create one or more new logical drives within the disk pool. A Free Capacity node can exist under each Disk Pool node. • Enhanced Flashcopy Images — This node represents a logical point-in-time image of a selected base logical drive. A base logical drive is a standard logical drive or a thin logical drive that is the source of an enhanced flashcopy image. • Enhanced Flashcopy Groups — This node represents the sequence of enhanced flashcopy images of the same base logical drive. • Enhanced Flashcopy logical drives — This node indicates that you have created a view of an enhanced flashcopy image. You create an enhanced flashcopy logical drive to allow a host to access an enhanced flashcopy image as if it were a logical drive. • Primary and Secondary logical drives — This node indicates whether the logical drive in the Enhanced Remote Mirroring logical drive pair is primary or secondary.

Table 5. Nodes on the Logical tab (continued)

Child nodes of the root node	Description of the child nodes
Arrays	<p>The Logical drive node and the Free Capacity node are standard child nodes. If the Enhanced Flashcopy Image premium feature is enabled, you can have the Enhanced Flashcopy image child nodes.</p> <ul style="list-style-type: none"> • Logical Drive — This node represents a configured and defined logical drive. An array supports standard logical drives only. Multiple logical drive nodes can exist under an Array node. • Free Capacity — This node represents a region of capacity that you can use to create one or more new logical drives within the array. Multiple Free Capacity nodes can exist under an array node. • Enhanced Flashcopy Images — This node represents a logical point-in-time image of a selected base logical drive. A base logical drive is a standard logical drive or a thin logical drive that is the source of an enhanced flashcopy image. • Enhanced Flashcopy Groups — This node represents the sequence of enhanced flashcopy images of the same base logical drive. • Enhanced Flashcopy Logical Drives — This node represents the enhanced flashcopy images of base logical drives that are visible to a host. • Primary and Secondary logical drives — This node indicates whether the logical drive in the Enhanced Remote Mirroring logical drive pair is primary or secondary. • Flashcopy Logical Drive — are child nodes of their associated base logical drives.
Consistency Groups	<p>If the Enhanced Flashcopy premium feature is enabled, you can have the following consistency group child nodes:</p> <ul style="list-style-type: none"> • Consistency Group — This node represents a grouping node which includes all the child nodes 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 image 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 member logical drives that are visible to a host.
Enhanced Global Mirror Groups	<p>These are special logical drives in the storage subsystem that are created as a resource for each controller in both local storage subsystems and remote storage subsystems. The controller stores duplicate information on the mirror repository logical drive, including information about remote writes that are not yet written to the secondary logical drives.</p>

Using the Host Mappings tab

Provides information about the two panes: the Host Mappings pane and the Defined Mappings pane in the Mappings tab on the Subsystem Management window.

The **Mappings** tab in the Subsystem Management window contains two panes: the Host Mappings pane and the Defined Mappings pane.

Host Mappings pane

The Host Mappings pane shows a tree-structured view of logical nodes that are related to storage partitions. Click the plus (+) sign or the minus (-) sign adjacent to a node to expand or collapse the view. You can right-click a node to open a pop-up menu that contains the applicable actions for that node.

The storage subsystem, or the root node, has these types of child nodes.

Table 6. Types of nodes in the Topology pane

Child nodes of the root node	Description of the child nodes
Undefined Mappings	The Undefined Mappings node has one type of child node: <ul style="list-style-type: none">• Individual Undefined Mapping: Represents a logical drive with an undefined mapping. Multiple Logical Drive nodes can exist under an Undefined Mappings node.
Default Group	Note: If the Storage Manager Storage Partitioning premium feature is disabled, all of the created logical drives are in the Default Group. A Default Group node has two types of child nodes: <ul style="list-style-type: none">• Host Group: Defined host groups that are not participating in specific mappings are listed. This node can have host child nodes, which can have child host port nodes.• Host: Defined hosts that are not part of a specific host group but are part of the Default Group and are not participating in specific mappings are listed. This node can have child host port nodes.
Unassociated Host Port Identifier	An Unassociated Host Port Identifier node has one type of child node. <ul style="list-style-type: none">• Host Port Identifier – Host port identifier that has not been associated with any host.
Host Group	A Host Group node has one type of child node. <ul style="list-style-type: none">• Host – Defined hosts that belong to this defined host group are listed. This node can have child host port nodes. Note: The host nodes that are child nodes of this host group can also participate in mappings specific to the individual host rather than the host group.
Host	A Host node has one type of child node: <ul style="list-style-type: none">• Host Ports: This node has child nodes that represent all of the host ports or single ports on a host adapter that are associated with this host.

The **Storage Partition** icon, when present in the Host Mappings pane, indicates that a storage partition has been defined for a host group, or a host. This icon also appears in the status bar if this feature has been enabled.

Defined Mappings pane

The Defined Mappings pane shows the mappings associated with a node that is selected in the Topology pane.

Table 7. Node information in the Defined Mappings pane

Column name	Description
Logical Drive name	The user-supplied logical drive name. The factory-configured access logical drive also appears in this column. Note: An access logical drive mapping is required for a storage subsystem with an in-band connection to enable the IBM Storage Manager to communicate with the storage subsystem. For a storage subsystem with out-of-band connections, you can remove an access logical drive mapping.
Accessible by	The Default Group, a defined host group, or a defined host that has been granted access to the logical drive in the mapping.
LUN	The LUN that is assigned to the specific logical drive that the host or hosts use to access the logical drive.
Logical Drive Capacity	The logical drive capacity in units of GB.
Type	The type of logical drive such as standard logical drive or flashcopy logical drive.

You can right-click a logical drive name in the Defined Mappings pane to open a menu. The menu contains options to change and remove the mappings.

Using the Hardware tab

Informs about the Hardware layout pane and Properties pane on the Hardware tab of the Subsystem Management window.

The **Hardware** tab contains two panes: the Hardware placement or Hardware layout pane on the left, and the Properties pane on the right. The Hardware placement pane provides a view of the hardware components in a storage subsystem, including their status.

The Hardware placement pane provides information for the hardware component that is selected in the Hardware pane. The information in the Properties pane is specific to each hardware component. If you select a controller icon in the Hardware pane, a list of properties for that controller is shown in the Properties pane. If you select a drive icon in the Hardware pane, a list of properties for that drive is shown in the Properties pane.

View

The **View Enclosure Components** command on each enclosure shows the status of the secondary components within the enclosure, such as power supplies, fans, and temperature sensors. You can select from **Drive type** and click **Show** in the

Hardware placement pane to identify drives of a particular type, speed, and capacity. A green triangle appears on top of the relevant drives.

Using the Setup tab

Lists the links and the tasks you can perform on the Setup tab of the Subsystem Management window.

The Subsystem Management window **Setup** tab provides links to the following tasks:

- Locate Storage subsystem
- Rename Storage subsystem
- Change Hardware View Order
- Set a Storage subsystem password
- Manage Premium features
- Create Storage
- Save Configuration

Optional Tasks:

- Manually Define Hosts
- Map Logical Drives
- Configure Ethernet Management Ports

Click a link to open the corresponding window.

Managing multiple software versions

Explains how the Storage Manager version is automatically selected for use in case the firmware version on multiple storage subsystems is different.

When you open the Subsystem Management window to manage a storage subsystem, the version of Storage Manager software that is applicable for the version of firmware opens. For example, you can manage two storage subsystems that use the Storage Manager software; one storage subsystem has firmware version 6.14, and the other has firmware version 7.5x. When you open a Subsystem Management window for one of the storage subsystems, the correct Subsystem Management window version is used. The storage subsystem with firmware version 6.14 uses version 9.14 of the Storage Manager software, and the storage subsystem with firmware version 7.5x uses version 10.5x of the Storage Manager software. You can verify the version that you are currently using by clicking **Help** > **About** in the Subsystem Management window.

Note: If you are managing multiple subsystems and these subsystems have different versions of controller firmware, the subsystem management windows may have a different look-and-feel depending on the controller firmware version.

Chapter 3. Installing Storage Manager

Lists the management-station operating systems for Storage Manager and contains links to various sections related to installing Storage Manager.

This chapter describes requirements and procedures for Storage Manager software installation.

For Storage Manager installation on unix-type operating systems, your system must have graphics capability to use the installation wizard. If your system does not have graphics capability, you can use the shell command to install Storage Manager without graphics. See “Installing Storage Manager with a console window in Linux, AIX, HP-UX, and Solaris” on page 31 for more information. You can also skip this section and install the stand-alone host software packages. Refer to “Installing Storage Manager manually” on page 32 and follow the procedures mentioned. All packages are included with the installation DVD. The supported management-station operating systems for Storage Manager are:

- AIX
- Windows 7, Windows Vista, Windows XP (Service Pack 2), Windows Server 2003, Windows Server 2008, and Windows Server 2012
- Linux: RHEL and SLES (x86, x86_64, and Linux on POWER)
- HP-UX (PA-RISC and IA64 editions)
- Oracle Solaris (SPARC and x86 editions)

Note: The IA64 version of the Storage Manager is not available for the IA64 edition of the Microsoft Windows and Linux operating systems. If you are managing a subsystem with controller firmware version earlier than 6.50.xx.xx, you must use an earlier version of Storage Manager software that is installed on a separate management workstation.

The installation instructions consist of the following sections:

- “Preinstallation requirements”
- “Installing the Storage Manager packages automatically with the installation wizard” on page 28
- “Installing Storage Manager packages manually” on page 32
- “Completing the Storage Manager installation” on page 34

To uninstall Storage Manager, see “Uninstalling Storage Manager” on page 33.

Attention: For cluster configurations, complete all applicable configuration procedures for each storage subsystem before you install the Storage Manager software on a second host or cluster server.

Preinstallation requirements

Informs you about the hardware, software, and configuration requirements of the management station

Note: With Storage Manager version 10.84.xx.xx or later, the minimum controller firmware must be 6.50.xx.xx or later. Controller firmware versions earlier than 6.50.xx.xx are not supported or managed.

The management station must also meet the following hardware, software, and configuration requirements:

- Microprocessor speed of 2 GHz or faster.
- Minimum of 2 GB of system memory. If any other applications are installed in the management station, additional memory might be required.
- Minimum of 1.5 GB of free disk space for the tool and for the saved support bundles.
- The TCP/IP stack must be enabled.

Installing the Storage Manager packages automatically with the installation wizard

Describes the installation process using the Storage Manager installation wizard on Windows and Unix-based operating systems.

You can install the Storage Manager software automatically by using the Storage Manager installation wizard, or you can install each package manually. This section describes the installation process for the Storage Manager installation wizard.

Before you install the Storage Manager software, read either of “Installing Storage Manager on Windows” or the “Installing Storage Manager on Linux, AIX, HP-UX, or Solaris” on page 29 section.

Installing Storage Manager on Windows

If your management station has a Windows operating system, complete the following steps to install Storage Manager with the installation wizard:

1. Download the files from the Storage Manager DVD, or from the System Storage Disk Support website, to a directory on your system. The default drive for Windows is C.
2. Double-click the **IBM DS Storage Manager package** (or SMIA) executable icon.
3. Follow the instructions in the Installation wizard to install the Storage Manager software. The default directory is

`C:\Program Files\IBM_DS`

or

`C:\Program Files(x86)\IBM_DS`

4. When you select the installation type, you can choose one of the following options:

Attention: Storage Manager SMIA package version 10.77.xx.xx and later will not install the MPIO DSM driver to support multipath in the host installation type or in the typical installation type when the SMIA package is installed in the server version of Microsoft Windows operating systems. There is a separate SMIA package for installing the MPIO DSM. The Storage Manager installer and the MPIO DSM installer are bundled in a single code package. This code package is available on the IBM support portal.

- **Typical (Full) Installation:** Installs Storage Manager software packages that are necessary for both managing the storage subsystem from the host and providing I/O connectivity to the storage subsystem
 - **Management Station:** Installs the packages that are required to manage and monitor the storage subsystem (SMclient)
 - **Host:** Installs the packages that are required to provide I/O connectivity to the storage subsystem (SMagent and SMutil)
 - **Custom:** Allows you to select which packages you want to install.
5. Install the MPIO DSM drive as required to support multipath by double-clicking the IBM DS Storage Manager MPIO DSM package and following the instructions in the installation wizard.

Note: This step applies to storage manager version 10.77.xx.xx and later only.

6. Click **Start > All Programs > DS Storage Manager 10 client > Storage Manager 10 client** to start the Storage Manager client program. Add the storage subsystems that you want to manage and monitor in the Enterprise Management window of the Storage Manager Client program.
7. Right-click the subsystem and select **Collect Support Data > Create/Edit Schedule**. The **Schedule Support Data Collection** window opens.
8. Set a schedule/time here and specify location to store support data during automatic data support collection for the storage subsystems you added.

Note: You can either schedule days of the month (1 to 31) or days of the week (Monday through Sunday) for automatic support data collection. A subsystem can have only one schedule. You can select the subsystems for which you want an identical schedule. Support data reports are saved on the local drive at the specified location. The report name includes the date on which the report was generated. The system allows a maximum of five support data reports. If a report is generated when five already exist, the oldest report is deleted and the new one is saved.

9. Select **Collect Support Data > Automatically**. The **Automatic Support Data Collection** window opens.

Note: For Automatic Support Data Collection, the Storage Manager monitor service must be enabled.

10. Select subsystems for automatic report in case of a critical event and specify location on your local drive to save it.

During the installation, the question **Automatically Start Monitor?** is displayed. This refers to the Microsoft Windows Event Monitor service. The Event Monitor must be enabled for both the automatic ESM synchronization and the automatic support bundle collection of critical events. To enable the Event Monitor, select **Automatically Start Monitor**.

To complete the Storage Manager installation, see “Completing the Storage Manager installation” on page 34.

Installing Storage Manager on Linux, AIX, HP-UX, or Solaris

If your management station has a Unix-based operating system, such as Linux, AIX, HP-UX, or Solaris, complete the following steps to install Storage Manager with the installation wizard:

1. Download the files from the Storage Manager DVD, or from the System Storage Disk Support website, to the root file system on your system.
2. Log in as root.
3. If the Storage Manager software package .bin file does not have executable permission, use the `chmod +x` command to make it executable.
4. Execute the .bin file and follow the instructions in the Installation wizard to install the software. The default directory is

/opt/IBM_DS

When you select the installation type, you can choose one of the following options:

- **Typical (Full) Installation:** Installs all Storage Manager software packages that are necessary for both managing the storage subsystem from this host and providing I/O connectivity to the storage
 - **Management Station:** Installs the packages that are required to manage and monitor the storage subsystem (SMruntime and SMclient)
 - **Host:** Installs the packages that are required to provide I/O connectivity to the storage subsystem (SMruntime, SMagent, and SMutil)
 - **Custom:** Allows you to select which packages you want to install.
5. Type SMclient in the console window and press Enter to start the Storage Manager Client program. Add the storage subsystems that you want to manage and monitor to the Enterprise Management window of the Storage Manager Client program.
 6. Right-click the subsystem and select **Collect Support Data > Create/Edit Schedule**. The **Schedule Support Data Collection** window opens.
 7. Set a schedule/time here and specify location to store support data during automatic data support collection for the storage subsystems you added.

Note: You can either schedule days of the month (1 to 31) or days of the week (Monday through Sunday) for automatic support data collection. A subsystem can have only one schedule. You can select the subsystems for which you want an identical schedule. Support data reports are saved on the local drive at the specified location. The report name includes the date on which the report was generated. The system allows a maximum of five support data reports. If a report is generated when five already exist, the oldest report is deleted and the new one is saved.

8. Right-click the subsystem. Select **Collect Support Data > Automatically**. The **Automatic Support Data Collection** window opens.

Note: For Automatic Support Data Collection, the Storage Manager monitor service must be enabled.

9. Select subsystems for automatic report in case of a critical event and specify location on your local drive to save it.
10. Configure, or Install and configure the multipath driver to manage paths to the mapped logical drive from the storage subsystem.

During the installation, the question **Automatically Start Monitor?** is displayed. This refers to the Event Monitor service. The Event Monitor must be enabled for both the automatic ESM synchronization and the automatic support bundle collection of critical events. To enable the Event Monitor, select **Automatically Start Monitor**.

To complete the Storage Manager installation, see “Completing the Storage Manager installation” on page 34.

Installing Storage Manager with a console window in Linux, AIX, HP-UX, and Solaris

Describes installation procedure for Storage Manager in a silent mode. This is applicable when Unix-based management stations do not have graphics adapter.

For a management station without a graphics adapter, the Storage Manager software package can be installed silently with the `-i silent` or `-i console` option. This installation method can also be used in a Windows operating-system environment.

The `-i silent` option causes the Storage Manager Software Installer package to be installed using the default installer settings. The `-i console` option prompts the user for installed options before the software installation starts, as the Installation wizard does. However, the prompts are displayed in console window text instead of graphical windows.

Portions of the Storage Manager console window installation text with the `-i silent` and `-i console` options are shown in the following example.

```
[usr@RHManaStation ~]# ./SMIA-LINUX-10.60.A5.17.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 environment...
```

```
Launching installer...
```

```
Preparing CONSOLE Mode Installation...
```

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

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

```
CHOOSE LOCALE BY NUMBER:
```

```
2
```

```
... ..
```

```
[usr@RHManaStation ~]# ./SMIA-LINUX-10.60.A5.17.bin -i silent
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 environment...
```

```
Launching installer...
```

```
Preparing SILENT Mode Installation...
```

```
=====
IBM System Storage DS Storage Manager 10
(created with InstallAnywhere by Macrovision)
-----
```

```

=====
Installing...
-----

[=====|=====|=====|=====]
[-----|-----|-----|-----]
... ..

```

Installing Storage Manager packages manually

For Unix-type operating systems such as AIX, Linux, Sun Solaris, and HP-UX, individual Storage Manager software packages are provided. See Table 8 for the installation sequence of each software package.

Use the procedure in this section to manually install the Storage Manager software on a management station. Be sure to install the packages in the correct order.

Important:

1. There is no manual installation option for Windows operating systems. For all installations of Storage Manager on Windows, the individual software packages are included in a single Storage Manager software installer.

Software installation sequence

Install the Storage Manager software packages in the sequence shown in Table 8.

Note: These packages are available for UNIX servers without a graphical user interface.

Table 8. Installation sequence of Storage Manager software packages

Step	Package
1	SMruntime
2	SMesm
3	SMclient ¹
4	SMagent
5	SMutil

¹SMclient is dependent on SMruntime, which is a Java compiler for SMclient. SMruntime must be installed first.

Installing Storage Manager manually

Before installing the Storage Manager software, make sure that the Storage Manager files are available in a directory on the system.

For your installation, modify the following commands as needed. No restart is required during the installation process. The verification process returns a table that describes the software installation, including the install package file name, version number, action, and action status.

1. Install the <SMpackage> by typing the command appropriate for your operating system.

Note: The manual install commands listed in the following table are *only* for UNIX-based operating systems.

Table 9. Examples of Storage Manager package install commands

Operating system	Package name	Install command
AIX	SMruntime.AIX-10.xx.xx.xx.bff	#installp -a -d /path_name/ SMruntime.AIX-10.xx.xx.xx.bff SMruntime.aix.rte
HP-UX	SMruntime_10.xx.xx.xx.depot	#swinstall -s /cdrom/HP-UX/ SMruntime_10.xx.xx.xx.depot
Solaris	SMruntime-SOL-10.xx.xx.xx.pkg	#pkgadd -d path/SMruntime-SOL- 10.xx.xx.xx.pkg
Linux on POWER	SMruntime-LINUX-10.xx.xx.xx-x.i586.rpm	#rpm -ihv SMruntime-LINUX- 10.xx.xx.xx-x.i586.rpm

2. Verify that the installation was successful by typing the command appropriate for your operating system.

Table 10. Storage Manager package installation verify commands

Operating system	Verify command
AIX	# ls1pp -ah <SMpackage>.aix.rte
HP-UX	# swverify -v <SMpackage>
Solaris	# pkginfo -l <SMpackage>
Linux on POWER	# rpm -qa grep <SMpackage>

If the verification process returns an error, contact your IBM service representative.

Uninstalling Storage Manager

Use the applicable procedure in this section to uninstall the Storage Manager, on a Windows or Unix-type operating system.

Uninstalling Storage Manager on a Windows operating system

To uninstall the software on a Windows operating system, complete the following steps:

1. Open the Control Panel window.
2. If you have Windows 2003 or Windows XP, double-click **Add/Remove Programs**. If you have Windows 2008 or Windows 2012, double-click **Program and Features**. The new window opens.
3. Select **IBM DS Storage Manager Host Software version 10.8x**, where *x* is the applicable version number of your software.
4. Click **Change/Remove** and follow the instructions in the Uninstall Storage Manager 10 wizard to uninstall the Storage Manager software. The process of uninstalling the software might leave files that were created by the Storage Manager after the installation was complete. These files might include trace files, repository files, and other administrative files. Delete these files manually to completely remove Storage Manager.

Uninstalling Storage Manager on a Linux, AIX, or Solaris operating system

To uninstall the software on a Unix-type operating system, complete the following steps:

1. Open the `/opt/IBM_DS/Uninstall` IBM System Storage DS Storage Manager 10 directory that contains the uninstaller binary.
2. Run the script `Uninstall_IBM_System_Storage_DS_Storage_Manager_10` script in the console window to uninstall Storage Manager software. The process of uninstalling the software might leave files that were not part of the original installation. These files might include trace files, repository files, and other administrative files. Delete these files manually to completely remove the Storage Manager.

Completing the Storage Manager installation

This section contains procedures for using the Enterprise Management and Subsystem Management features of the Storage Manager to complete the Storage Manager installation tasks for all host operating systems.

To complete a Storage Manager installation, the following procedures must be performed:

- Perform an initial automatic discovery of storage subsystems
- Perform an initial manual discovery of storage subsystems
- Name the storage subsystems
- Set up alert notifications
- Create a schedule to automatically collect support data for the managed storage subsystems
- Enable collect support data in case of a critical event
- Configure iSCSI settings for storage subsystems with iSCSI ports
- Verify and upgrade the Controller code, ESM, and disk drive firmware to the latest version as prescribed on the IBM support portal
- Enable Storage subsystem premium features
- Save a storage subsystem profile and support data

Each of these procedures is described in detail in the following sections.

The Enterprise Management window opens when you start the 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
- Perform batch storage subsystem management tasks with the Script Editor

Performing an automatic discovery of storage subsystems

Complete the following steps to perform an initial automatic discovery of storage subsystems:

1. For Windows operating system, click **Start > All Programs > DS Storage Manager 10 Client > DS Storage Manager 10 Client**. For UNIX-type operating system, open the console window. Type `SMclient` and press Enter. The Storage Manager client software starts and displays the Enterprise Management window and the Confirm Initial Automatic Discovery window.

2. Click **Yes** to begin an initial automatic discovery of hosts and storage subsystems that are attached to the local subnetwork.

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

Note: The Enterprise Management window can take one minute or more to refresh after an initial automatic discovery. If the storage subsystem is not discovered automatically, check the network connections (out-of-band management) or the server HBA port to storage subsystem host port connections. To try to add the subsystem manually, click **Edit > Add Storage Subsystem**.

3. Verify that each host and storage subsystem is displayed in the Enterprise Management window.
 - If a host or storage subsystem is not displayed, complete the following tasks:
 - a. Check the hardware and hardware connections for possible problems. See the *Installation, User's, and Maintenance Guide* for your storage subsystem for specific procedures.
 - b. See the Enterprise Management online help for additional information about discovering storage subsystems.
 - c. 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 and the gateway information is defined for the Ethernet ports. For more information about the storage subsystem ethernet management ports, refer to "Setting up IP addresses for storage subsystem controllers" on page 7. If you are using the host-agent method (commonly known as in-band management), make sure that the Fibre Channel, SAS, or iSCSI connection between the host and storage subsystems is made.
 - d. 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, a partially-managed device error message is displayed when you attempt 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 storage subsystem.

- If you are using the host-agent-management method, complete the following steps:
 - a. Make sure that the SMagent is installed in the host.
 - b. Verify that you have a Fibre Channel, SAS, or iSCSI connection from the storage subsystems to the host on which the SMagent installed. Check the SAN switch zoning or VLAN configuration as required.
 - c. Verify that all of the preparation steps are complete, and then perform the following steps:
 - 1) Run the hot_add utility.
 - 2) Restart the SMagent.
 - 3) Right-click the host, and click **Tools > Rescan Hosts** in the Enterprise Management window.

Note: In certain situations, a storage subsystem might be duplicated in the **Device** tab tree view 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.

4. 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. Verify that the storage subsystem is powered on and complete start-of-day process. Then use the **Add Device** option to add it to the management domain again. See the Enterprise Management window online help for instructions for removing and adding devices.

Performing a manual discovery of storage subsystems

You can add hosts or storage subsystems manually; use this option to selectively manage a group of storage subsystems from an SMclient. You can also use this option to add devices to be managed that were not discovered during the SMclient initial discovery.

In the Enterprise Management window, click **Edit > Add Storage Subsystem**. The **Add New Storage Subsystem - Manual** window opens. You can add Storage subsystems for management. Refer to the online help to know how to add subsystems.

Important:

1. 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.
2. 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.

Setting a storage subsystem management password

The storage subsystem management password functionality differs between certain combinations of Storage Manager and controller firmware versions.

You are prompted with a window to set the subsystem management password every time you start the Subsystem Management window of the storage subsystem that did not have the password set. In addition, the password times out after a certain duration of Subsystem Management window inactivity. The password must be between 8 and 30 characters and contain at least one number, one lower-case letter, one upper-case letter, and one non-alphanumeric character (for example, < > ! @ + #). Spaces are not permitted, and it is case-sensitive.. Storage subsystems with controller firmware version 7.70.xx.xx do not allow the subsystem management window to be opened if the subsystem management password is not set. There are no such restrictions for other controller firmware versions.

Important: Make sure that the password information is kept in a safe and accessible place. Contact IBM technical support for help if you forget the password to the storage subsystem.

Naming storage subsystems

As you set up your network, decide on a naming convention for the storage subsystems. For example, you might use the device type followed by the serial number: 1815 1312345 XXXX xxx xxxx. When you install the Storage Manager software and start it for the first time, all storage subsystems in the management domain are shown as <unnamed>. Use the Subsystem Management window to rename each storage subsystem.

Consider the following factors when you name storage subsystems:

- There is a 30-character limit. All leading and trailing spaces are deleted from the name.
- Use a unique, meaningful naming scheme that is easy to understand and remember.
- Avoid arbitrary names or names that might quickly lose their meaning.
- The software adds the prefix Storage Subsystem when it displays storage subsystem names. For example, if you name a storage subsystem Engineering, it is displayed as Storage Subsystem Engineering.

To name your storage subsystems, complete the following steps:

1. In the Enterprise Management window, right-click the storage subsystem and select **Rename**. The Rename Storage Subsystem window opens.

Note: If any of the hosts are running path failover drivers, update the storage subsystem name in your path failover driver configuration file before you reboot the host system to establish uninterrupted access to the storage subsystem.

2. Type the name of the storage subsystem and click **OK**.
3. Click **Yes** on the warning screen.
4. Repeat this procedure for each unnamed storage subsystem. For more information, see the topic about renaming storage subsystems in the Subsystem Management window online help.
5. Proceed to “Setting alert notifications.”

Setting alert notifications

After you add devices to the management domain, you can set alert notifications to report critical events on the storage subsystems. The following alert-notification options are available:

- Notification to a designated network-management station using Simple Network Management Protocol (SNMP) traps
- Notification to designated email addresses
- Notification to designated alphanumeric pagers (requires separately supplied software to convert email messages)

Note: You can monitor storage subsystems only within the management domain. If you do not install the Event Monitor service, the Enterprise Management window must remain open. If you close the window, you will not receive any alert notifications from the managed storage subsystems. See the Enterprise Management window online help for additional information.

Alert notification with SNMP traps

To set up alert notification to a network-management station using SNMP traps, complete the following steps:

1. Insert the Storage Manager DVD into the DVD drive on a network-management station. You must set up the designated management station only once.
2. Copy the SMxx.x.MIB file from the SMxxMIB directory to the network-management station.
3. Follow the steps that are required by your network-management station to compile the management information base (MIB) file. (For details, contact your network administrator or see the documentation for your particular storage management product.)

Alert notification without SNMP traps

To set up alert notification without using SNMP traps, click **Storage subsystem > Edit > Configure alerts** on the Enterprise Management window.

Automatic support bundle collection

Starting with Storage Manager Version 10.83, the ability to automatically collecting support data from the managed storage subsystems periodically is built into the storage manage client software. The information in the saved support data collections might be helpful in troubleshooting and recovery of the storage subsystem in the event of a catastrophic failure. To create a schedule for automatically collecting the support data, select **Tools > Collect Support Data > Create/Edit** in the Enterprise management window. When the **Schedule Support Data Collection** window opens, select the storage subsystem and click **Create/Edit** to create the schedule to collect support data daily, weekly, monthly or yearly. You can create the schedule for multiple subsystems at the same time by holding down the Control key while selecting the subsystems. However, IBM recommends staggering the data collection time when managing multiple subsystems. Refer to the online help for more information about collecting support data periodically. In addition to periodical automatic support data collection, the storage manager client also has the ability to collect support data when a critical event occurs in the managed storage subsystem. Select **Tools > Collect Support Data > Automatically** in the Enterprise management window to configure this. Make changes as required when the **Automatic Support Data Collection** window opens. Refer to the online help for more information about it.

The schedule is stored in the management station where it is defined. The management station must be up and running and must have management connection to the storage subsystems for the support bundle to be automatically created. The filenames of the saved files also include the date. The storage manager code maintains a maximum of five files. If five files already exist, the code deletes the oldest file before saving the new one. If there is not enough disk space to save a new file, the storage manager code deletes files beginning with the oldest, until enough space is created.

Note: The Storage Manager client need not be running for the support data to be collected. However, the IBM DS Storage Manager support monitor service must be running. Also, to prevent deletion of support data when the repository directory is full, select repository location with adequate space.

Using SM scripts for configuring and managing automatic support bundle collection

Instead of using the menu options in the Enterprise management window, you can configure automatic support bundle collection, display current schedule, and schedule collection from the command line interface of your management station. To know more about these commands, see the *IBM System Storage DS3000, DS4000, and DS5000 Command Line Interface and Script Commands Programming Guide* or refer to the online help. You must run these commands from the management station on which you installed the Storage Manager and not from the script window that is launched from the Storage Manager Enterprise management window. In addition, do not specify the storage subsystem IP address as part of the SMcli command and do not precede the supportBundle command with '-c' parameter. Use the subsystemName of the managed storage subsystem as script command value. To know the subsystemName, run the following smcli command with -d parameter.

```
C:\Program Files (x86)\IBM_DS\client>smcli -d
DS3400A          fe80:0:0:0:2a0:b8ff:fe5a:ae42
DS3524-DTL139140 ds3524dta.net.com ds3524dtb.net.com
DS3524-DTL       ds3524dt1.net.com ds3524dt2.net.com
DS5100-DTL       ds5k-a.net.com ds5k-b.net.com
```

DS3400A, DS3524-DTL139140, DS3524-DTL, and DS5100-DTL are the names of the storage subsystems. To enable a schedule to automatically collect support bundles on Sunday and Tuesday each week at 2:00AM on the storage subsystem DS3524-DTL139140

```
C:\Program Files (x86)\IBM_DS\client>smcli -supportBundle schedule
enable DS3524-DTL139140 daysofweek=[Sunday Tuesday] startTime=02:00
```

Both these commands are run from the command line interface of the management station.

Note: If your management station is on Windows operating system, you must be an administrator to run these commands.

Managing iSCSI settings

Click the **Setup** tab in the Subsystem Management window.

Note: The link to **iSCSI Manage settings** or **Configure iSCSI Host Ports** on the Subsystem Management window is available only for storage subsystems that support an iSCSI host attachment. As of the date of this document, the following storage subsystems support iSCSI host attachment:

- DS3300
- DS3500
- DCS3700 with Performance Module Controllers
- DS3950
- DS5020
- DS5100/5300

The following iSCSI options are available from the **Storage Subsystem management** menu and are described in the following sections:

Note: The menu selection in these iSCSI options changes according to the version of the controller firmware. Refer to the online help for the appropriate menu options.

- “Changing target authentication”
- “Entering mutual authentication permissions”
- “Changing target identification”
- “Changing target discovery”
- “Configuring iSCSI host ports”
- “Viewing or ending an iSCSI session”
- “Viewing iSCSI statistics” on page 41

Changing target authentication

Select **Change Target Authentication** to specify the target Challenge Handshake Authentication Protocol (CHAP) secret that the initiator must use during the security negotiation phase of the iSCSI login. By default, **None** is selected. To change the selection, click **CHAP**, and then enter the CHAP secret. You can also select the option to generate a random secret. This enables 1-way CHAP.

Entering mutual authentication permissions

Before you select **Enter Mutual Authentication Permissions**, you must define a host port for the initiator and enable Target Authentication. After the host port is listed, select the host from the list and click **Chap Secret** to specify the secret that is passed to the initiator from the target and authenticate it. This enables Mutual CHAP (2-way).

Changing target identification

Select **Change Target Identification** to specify a target alias that is to be used during device discovery. You must provide a unique name that consists of fewer than 30 characters for the target.

Note: You will connect to the target with the fully qualified IQN that is listed above the alias.

Changing target discovery

Select **Change Target Discovery** to perform device discovery with the iSCSI simple naming service (iSNS). After you select this option, select the **Use iSNS Server** check box. You can also select whether the iSNS server is discovered using a DHCP server on your network, and you can manually specify an Internet Protocol version 4 (IPv4) or IPv6 address. When you click the **Advanced** tab, you can assign a different TCP/IP port for your iSNS server for additional security.

Note: To provide the required port login information for correct device discovery, all iSCSI ports must be able to communicate with the same iSNS server.

Configuring iSCSI host ports

Select **Configure iSCSI Host Ports** to configure all of the TCP/IP settings. You can choose to enable or disable IPv4 and IPv6 on all of the ports. You can also statically assign IP addresses or let them be discovered using DHCP. Under **Advanced IPv4 Settings**, you can assign VLAN tags (802.1Q) or set the Ethernet priority (802.1P). Under **Advanced Host Port Settings**, you can specify a unique iSCSI TCP/IP port for that target port. You can also enable jumbo frames from this option. The supported frame sizes are 1500 and 9000.

Viewing or ending an iSCSI session

Select **View/End iSCSI Sessions** to view all of the connected iSCSI sessions to the target. From this page, you can also close an existing session by forcing a target ASYNC logout of the initiator session.

Viewing iSCSI statistics

Select **View iSCSI Statistics** to view a list of all iSCSI session data, for example, the number of header digest errors, the number of data digest errors, and successful protocol data unit counts. You can also set a baseline count after a corrective action to determine whether the problem is solved.

Using an iSNS server

There are many considerations for using an iSNS server correctly. Be sure to correctly assign your iSNS server address that is provided during the DHCP lease discovery of your initiator or target. This enables ease of discovery when you use initiator-based solutions. If you are unable to do this and you must manually assign the iSNS server to your software or hardware initiators, you must make sure that all of the storage subsystem iSCSI ports and iSCSI initiators are in the same network segment (or make sure that the routing between the separate network segments is correct). If you do not do this, you will be unable to discover all ports during the iSCSI discovery process, and you might not be able to correctly perform a controller or path failover.

Using DHCP

Do not use DHCP for the target portals. If you use DHCP, you must assign DHCP reservations so that leases are maintained consistently across restarts of the storage subsystem. If static IP reservations are not provided, the initiator ports can lose communication to the controller and might not be able to reconnect to the device.

Using supported hardware initiators

Before you install and configure these adapters, ensure that you have installed the latest management application and firmware. Then, configure each adapter one at a time. In a configuration with a single-controller subsystem, in which iSCSI adapters and target ports are in the same network segment, each adapter can connect to any target port. In a complex configuration, each adapter can connect to a controller device. To ensure that failover is performed correctly, connect each iSCSI adapter in the server in one of the following:

- **- Single iSCSI adapter in the server** - The iSCSI adapter port must be able to login both controller A and B iSCSI host ports. The iSCSI port is configured as multihome with controller A and B iSCSI port subnets.
- **- Multiple iSCSI adapters in the server** - Each adapter port is allowed a single path to each controller iSCSI host ports. Each iSCSI adapter port and its associated iSCSI controller host port must be on separate subnet from other pairs of iSCSI adapter and controller host ports

If you use a Qlogic hardware initiator adapter, complete the following steps to log in to all the available target ports from the hardware initiator. For other hardware initiator adapters, refer to the publications of those hardware initiator adapters for information on logging in to all the available target ports.

For the list of supported hardware initiators, go to <http://www.ibm.com/systems/support/storage/config/ssic>.

To log in to all available target ports from the Qlogic hardware initiator, complete the following steps.

Note: Failure to perform the steps in the following procedure might result in path failover inconsistencies and incorrect operation of the storage subsystem.

1. Start the SANsurfer management utility.
2. Connect to the system that is running the qlremote agent.

3. Select the adapter that you want to configure.
4. Select either **Port 0** or **Port 1** for the adapter.
5. Click **Target Settings**.
6. Click the plus sign (+) in the far right of the window.
7. Type the IPv4 or IPv6 address of the target port to which you want to connect.
8. Click **OK**.
9. Select **Config Parameters**.
10. Scroll until you see ISID. For connection 0, the last character that is listed must be 0. For connection 1, it must be 1, for connection 2, it must be 2, and so on.
11. Repeat steps 6 through 10 for each connection to the target that you want to create.
12. After all of the sessions are connected, select **Save Target Settings**. If you are using the QLogic iSCSI Single-Port or Dual-Port PCIe HBA for IBM System x to support IPv6, you must allow the host bus adapter firmware to assign the local link address.

Using IPv6

The storage subsystem iSCSI ports support the Internet Protocol version 6 (IPv6) TCP/IP. Note that only the final four octets can be configured if you are manually assigning the local link address. The leading four octets are fe80:0:0:0. The full IPv6 address is required when you are attempting to connect to the target from an initiator. If you do not provide the full IPv6 address, the initiator might fail to be connected.

Configuring network settings for iSCSI host attachment

If you use a storage subsystem that supports iSCSI host attachment in a complex network topology, you must address a few challenges. If possible, isolate the iSCSI traffic to a dedicated network. If this is not possible and you are using a hardware-based initiator, the Keep Alive timeout must be set to 120 seconds. To set the Keep Alive timeout, complete the following steps:

1. Start the SANsurfer Management Utility and connect to the server.
2. Select the adapter and the adapter port that is to be configured.
3. Select the port options and firmware.

The default connection timeout is 60 seconds. This setting is satisfactory for simple network topologies. However, in a more complex configuration, if a network convergence occurs and you are not using Fast Spanning Tree and separate spanning tree domains, you might incur I/O timeouts. If you are using a Linux iSCSI software initiator, modify the ConnFailTimeout parameter to account for the spanning tree issue. The ConnFailTimeout value must be set to 120 seconds.

Configuring Maximum Transmission Unit settings

All devices on a link that are expected to communicate with each other (such as devices on the same VLAN) must be configured with the same Maximum Transmission Unit (MTU) size. The MTU size is a configuration item, or it is hard-coded in the device, and it is not negotiated between endpoints during login or connection establishment. If a device receives a packet that is larger than the MTU size, it drops the packet. If a router receives a packet whose size does not exceed the MTU size of the link on which it was received but exceeds the MTU size of the forwarding link, the router either fragments the packet (IPv4) or returns a packet too large ICMP error message. Make sure that all of the components on a network link are using the same MTU size value.

For storage subsystems that support iSCSI, the default MTU setting is 1500 bytes. There is an option to select 9000 bytes for jumbo frames. For end-to-end jumbo frames to work effectively, jumbo frames (large MTU) must be enabled on all components (host, switch, routers, and targets). If jumbo frames are not enabled on all components, one or more of the following conditions might occur:

- Frames are dropped.
- No connections are dropped because of error messages about the packet being too large.
- Jumbo frames are fragmented.

Microsoft iSCSI Software Initiator considerations

The native multipath I/O (MPIO) that is provided with the Microsoft iSCSI Software Initiator (version 2.03 or later) is not supported. You must use the DSM that is provided with Storage Manager to make sure that failover and I/O access are correct. If the native MPIO from the Microsoft iSCSI Software Initiator is used, it causes unwanted effects.

Downloading controller firmware, NVSRAM, ESM firmware

This section provides instructions for downloading storage subsystem controller firmware, NVSRAM, storage enclosure ESM firmware, and drive firmware. Normally, the storage subsystem firmware download sequence is as follows:

1. Controller firmware
2. Controller NVSRAM
3. ESM firmware
4. Drive firmware

Review the readme file that is provided with updated controller firmware, NVSRAM, ESM firmware, and drive firmware for any necessary changes to the firmware download sequence.

Important:

1. The following procedures assume that you are using the latest controller firmware version. Access the latest versions of storage subsystem controller firmware, NVSRAM, and storage enclosure ESM firmware on the IBM Support Portal at <http://www.ibm.com/support/entry/portal>. For the most recent Storage Manager readme files for your operating system, see “Finding Storage Manager software, controller firmware, and readme files” on page xiv.
2. IBM supports the storage subsystem controller and ESM firmware download with I/O, sometimes called *concurrent firmware download*, with some storage subsystems. Before you proceed with concurrent firmware download, review the readme file that is packaged with the firmware code or your operating-system Storage Manager host software for any restrictions.
3. Suspend all I/O activity while you download firmware and NVSRAM to a storage subsystem with a single controller. If you do not suspend I/O activity, the host server will have failed I/O requests because you have redundant controller connections between the host server and the storage subsystem.
4. Always check the storage subsystem controller firmware readme file for any controller firmware dependencies and prerequisites before you apply the firmware updates to the storage subsystem. Updating any components of the storage subsystem firmware without complying with the dependencies and prerequisites might cause downtime (to fix the problems or recover).

5. Downgrading the controller firmware is not a supported function. This option should be used **only** under the direction of IBM Support. Downgrading from 07.xx to 06.xx firmware levels is not supported and will return an error if attempted.

If your existing controller firmware is 06.1x.xx.xx or later, you have the option to select the NVSRAM for download at the same time that you upgrade or download the new controller firmware. Additionally, you have the option to download the firmware and NVSRAM immediately but activate it later, when it might be more convenient. See the online help for more information.

Note: The option to activate firmware at a later time is not supported on the DS4400 storage subsystem.

Determining firmware levels

Before you download a firmware upgrade, ensure that you know the current firmware version. There are three different methods to determine storage subsystem, storage enclosure, drive, and ESM firmware versions. Each method uses the Storage Manager client that manages the storage subsystem with the attached storage enclosure.

Method one:

Go to the Subsystem Management window and select the menu option to display the storage subsystem profile. When the Storage Subsystem Profile window opens, click the **Summary** tab and scroll through the **Monitor** page to locate the following information. The **View Firmware Inventory or View Storage Subsystem Profile** page contains all of the profile information for the entire storage subsystem, including the firmware version numbers. An example is as under -

Storage subsystem controller firmware version

View Storage Subsystem Profile:

See the following example of profile information.

```
FIRMWARE INVENTORY
IBM DS Storage Manager 10
  SMW Version:      10.84.G5.21
  Report Date:      Tue Oct 09 21:13:34 CST 2012

Storage Subsystem
  Storage Subsystem Name:  DCS3700
  Current Package Version: 07.84.39.00
  Current NVSRAM Version:  N1818D37R0784V04
  Staged Package Version:  None
  Staged NVSRAM Version:  None
```

Controllers

```
Location:      Enclosure 1, Slot A
Current Package Version: 07.84.39.00
Current NVSRAM Version:  N1818D37R0784V04
Board ID:      2660
Sub-Model ID:  162

Location:      Enclosure 1, Slot B
Current Package Version: 07.84.39.00
Current NVSRAM Version:  N1818D37R0784V04
Board ID:      2660
Sub-Model ID:  162
```

Power Supplies

Location:	Enclosure 1 Right
Firmware Version:	Not Available
Location:	Enclosure 1 UNKNOWN
Firmware Version:	Not Available

Drive

Enclosure, Drawer, Slot:	Manufacturer:
Product ID:	Drive Type:
Capacity:	
Drive Firmware Version:	FPGA Version: (SSD only)
Enclosure 1, Drawer 1, Slot 2	IBM-ESXS
ST9300603SS F	Serial Attached SCSI (SAS)
278.896 GB	
B53B	Not Available
Enclosure 1, Drawer 1, Slot 3	IBM-ESXS
CBRCA300C3ETS0 N	Serial Attached SCSI (SAS)
278.896 GB	
C610	Not Available

Method two:

Complete the applicable procedure from the following options to obtain the specified firmware version.

To obtain the controller firmware version:

Click the **Controller** icon in the **Hardware** tab of the Subsystem Management window. Controller information displays in a new window.

You must perform this action for each controller.

To obtain the drive firmware version:

Click the **Controller** icon in the **Hardware** tab of the Subsystem Management window. Drive firmware information displays in a new window.

You must perform this action for each controller.

To obtain the ESM firmware version:

Click the **Controller** icon in the **Hardware** tab of the Subsystem Management window. ESM firmware information displays in a new window.

You must perform this action for each controller.

Downloading controller and NVSRAM firmware

Note: Perform a Collect All Support Data operation before you upgrade the controller firmware and NVSRAM. See “Critical event problem solving” on page 245 for the data-collection procedures.

This section provides instructions for downloading storage subsystem controller firmware and NVSRAM. Normally, the storage subsystem firmware download sequence starts with controller firmware, followed by the NVSRAM, the ESM firmware, and the drive firmware.

Important: If you are upgrading from 06.xx to 07.xx, you must use the Controller Firmware Upgrade Tool. See “Using the IBM System Storage Controller Firmware Upgrade Tool” on page 46. If you already have the 07.xx firmware level, you are not required to use the Controller Firmware Upgrade Tool to upgrade to another 07.xx level. However, the Upgrade Tool has diagnostic capabilities that might be beneficial.

To download firmware version 06.1x.xx.xx or later, and NVSRAM, complete the following steps:

1. From the Enterprise Management window, select a storage subsystem.
2. Click **Tools > Upgrade Controller Firmware**. The Upgrade Controller Firmware window opens.

Note: If the controller firmware is 7.77.xx.xx or later, the system automatically runs a pre-upgrade check, which takes several minutes. The controller firmware upgrade proceeds only if the pre-upgrade check is satisfactory. With controller firmware versions 06.1x.xx.xx and later installed on the storage subsystem, you can download the NVSRAM file with the firmware file. This download feature is not supported in storage subsystems with controller firmware 05.4x.xx.xx or earlier. If the existing controller firmware is version 05.4x.xx.xx or earlier, only a window for downloading firmware is displayed.

The Storage Manager software checks the status of each storage subsystem and lists storage subsystems that need an update.

3. Select all storage subsystems that you want to upgrade. Click **Firmware**. The Download Firmware window opens.

Note: If you want to upgrade multiple subsystems at a time, all those subsystems must be of the same type.

Using the IBM System Storage Controller Firmware Upgrade Tool

The following information describes how to use the IBM System Storage Controller Firmware Upgrade Tool to upgrade your DS4800, DS4700 or DS4200 Express® controller firmware from 06.xx to 07.xx.

Important:

- Use the Controller Firmware Upgrade Tool only when you are migrating DS4800, DS4700, and DS4200 controllers from version 06.xx to version 07.xx.
- This tool is not intended for standard upgrades for controller, ESM, or drive firmware (upgrades within the 6.xx.xx.xx code thread or 7.xx.xx.xx code thread). To perform a standard controller, ESM, or drive-firmware upgrade, see “Downloading controller firmware, NVSRAM, ESM firmware” on page 43.
- With Storage Manager version 10.50 and later, the Controller Firmware Upgrade Tool has become part of the Enterprise Management window and is no longer a separate tool. Click **Tools > Upgrade Controller Firmware** to open a Controller Firmware Upgrade window.

Attention: Make sure that the firmware that you download is compatible with the Storage Manager software that is installed on your storage subsystem. If incompatible firmware is downloaded, you might lose access to the drives in the storage subsystem. If necessary, upgrade Storage Manager before you use the IBM System Storage Controller Firmware Upgrade Tool.

Before you use the IBM System Storage Controller Firmware Upgrade Tool, make sure that you understand the following guidelines.

- All data for DS4800, DS4700, and DS4200s storage subsystems must be backed up, and existing system configurations must be saved. After the tool has completed an upgrade, controllers cannot return to previous firmware version levels.
- You must perform the upgrade offline. You must perform the overall installation of Storage Manager into an existing host environment online. For most failover drivers to take effect, you must reboot the host.

- All devices must have Optimal status before you download firmware. Use the Healthcheck utility to verify Optimal status on the devices.
- You must also check the current firmware level.
- Do not make changes to your configuration or remove drives or enclosures during the upgrade process.
- For information about the current firmware versions, see “Finding Storage Manager software, controller firmware, and readme files” on page xiv to find out how to access the most recent Storage Manager readme files on the web.

Checking the device health conditions: To determine the health condition of your device, complete the following steps:

1. From the Subsystem Management window in the Storage Manager, right-click the storage subsystem. Storage Manager establishes communication with each managed device and determines the current device status.

There are six possible status conditions:

- **Optimal:** Every component in the managed device is in optimal working condition.
 - **Needs Attention:** There is a problem with the managed device that requires intervention to correct it.
 - **Fixing:** A Needs Attention condition has been corrected, and the managed device is currently changing to an Optimal status.
 - **Unresponsive:** The management station cannot communicate with the device or with one controller or both controllers in the storage subsystem.
 - **Contacting Device:** Storage Manager is establishing contact with the device.
 - **Needs Upgrade:** The storage subsystem is running a level of firmware that is no longer supported by Storage Manager.
2. If the Status is not Optimal, click **Recovery Guru** in the Subsystem Management window for detailed information about system recovery before firmware download. If you fail in system recovery, contact an IBM technical support representative for fault resolution.

Opening and using the Controller Firmware Upgrade Tool: To use the Controller Firmware Upgrade Tool, click **Tools > Firmware Upgrade** in the Enterprise Management window. The Firmware Upgrade window opens. The Firmware Upgrade Tool automatically completes a diagnostic check on these subsystems to determine if they are healthy enough to perform a controller firmware upgrade.

Note:

- For any condition other than Optimal, you must contact IBM support for assistance. See “Software service and support” on page xvii for additional information.
- You can only upgrade from a major release to a major release (for example, 06.xx. to 07.xx) with this tool. Do not attempt to perform this type of firmware upgrade in the Subsystem Management window.
- After you have upgraded to the 07.xx firmware level, you do not need to use the firmware upgrade tool. Use the Subsystem Management window to perform any future firmware upgrades.

For more information about using the tool, click the **Help** button in the Controller Firmware Upgrade Tool.

Adding a storage subsystem: To add a storage subsystem by using the Controller Firmware Upgrade Tool, complete the following steps:

1. Click **Add**. The Select Addition Method window opens.
2. Click **Automatic** or **Manual**.
3. Click **OK** to begin adding storage subsystems.
4. If you want to see issues with an added storage subsystem that might impede firmware upgrading, click **View Log**.

Downloading the firmware:

1. Select the storage subsystem that you want to activate. The **Download** button is enabled.
2. From the Enterprise Management window toolbar, click **Tools > Upgrade Firmware**. The Download Firmware window opens.
3. Click **Browse** and select the controller firmware file that you want to download from its directory.
4. Click **Browse** and select the NVSRAM file from its directory.
5. Click **OK**. The firmware download begins. A status bar is displayed in the Controller Firmware Upgrade window.

Viewing the IBM System Storage Controller Firmware Upgrade Tool log file:

The Controller Firmware Upgrade Tool log file documents any issues with the storage system that might prevent you from updating the firmware. If you encounter any problems when you upgrade the firmware, click **View Log** to open the log file. Correct the issues in the log file before you try to download the firmware again.

Downloading ESM firmware

This section provides instructions for downloading storage enclosure ESM firmware. Normally, the storage subsystem firmware download sequence starts with the controller firmware, followed by the NVSRAM, ESM firmware, and drive firmware.

To download the ESM firmware, complete the following steps:

1. In the System Management window, select **Upgrade > ESM firmware**. The Download Environmental Card Firmware window opens.
2. Click **Select All** to direct the download to all storage enclosures. You can also select one storage enclosure, or you can select multiple enclosures by pressing Ctrl while you select the enclosures.

Note: If you have selected multiple enclosures, suspend all I/O activity while the ESM firmware downloads. If select only one storage enclosure at a time, you can download ESM firmware while the server conducts I/O activity.

3. Click **Browse** to identify and select the file name of the ESM firmware file, and click **Start** to begin the ESM firmware download.
4. In the Confirm Download window, type yes and click **OK** to start the download process.
5. After the ESM firmware download to all selected enclosures is complete, click **Cancel** to close the window.

When you install a new ESM into an existing storage enclosure in a storage subsystem that supports automatic ESM firmware synchronization, the firmware in the new ESM is synchronized automatically with the firmware in the existing ESM. This resolves any ESM firmware mismatch conditions automatically.

To enable automatic ESM firmware synchronization, make sure that your system meets the following requirements:

- The Storage Manager Event Monitor must be installed and running.
- The storage subsystem must be defined in the Enterprise Management window of the Storage Manager client (SMclient).

Note: As of the date of this document, the Storage Manager supports automatic ESM firmware synchronization with EXP710 and EXP810 storage enclosures only. Contact IBM for information about support for other types of storage enclosures in the future. To correct ESM firmware mismatch conditions in storage enclosures without automatic ESM firmware synchronization support, you must download the correct ESM firmware file with the ESM firmware download menu function in the SMclient Subsystem Management window.

Downloading drive firmware

This section provides instructions for downloading drive firmware. Up to four different drive types can have drive firmware updated at the same time. Drives are considered as different drive types if they report different product IDs when an inquiry is made. See the online help for additional information.

Important:

1. The following procedures assume that you have the latest controller firmware version. If you have an earlier firmware version, see “Finding Storage Manager software, controller firmware, and readme files” on page xiv to obtain the applicable firmware version documentation.
2. IBM supports firmware download with I/O, sometimes referred to as *concurrent firmware download*. This feature is not supported for drive firmware. You must schedule downtime for drive and ATA translator firmware upgrade.

To download drive firmware for Storage Manager, complete the following steps:

1. Before you start the drive-firmware download process, complete the following tasks:
 - a. Stop all I/O activity before you download drive firmware to a storage subsystem.
 - b. Complete a full backup of all data on the drives that you select for the firmware upgrade.
 - c. Unmount the file systems on all logical drives that access the drives that you select for the firmware upgrade.
2. From the Enterprise Management window, select a storage subsystem.
3. On the Subsystem Management window menu bar, click **Upgrade > Drive firmware**. The Introduction page opens. Read the instructions and click **Next**.

Note: Using Storage Manager, you can simultaneously download and update up to four different firmware packages.

4. Click **Add** to locate the server directory that contains the firmware that you plan to download.
5. Select the firmware file that you plan to download and click **OK**. The file is listed in the Selected Packages window.
6. Repeat steps 4 and 5 for up to four packages that you plan to download firmware and click **Next**. Additional files are listed in the Selected Packages window.
7. After you specify the firmware packages for download, click **Next**.

8. The **Compatible Drives** page lists drives that are compatible with the firmware package types that you selected. From that list, select the drives to which you plan to download the drive firmware. You can press and hold the Ctrl key while you select multiple drives individually, or you can press and hold the Shift key while you select multiple drives that are listed in series. Click **Select All** to select all drives.

Note: The firmware that you plan to download must be listed on the Compatible Drives page. If the product ID of your drives matches the firmware type and it is not listed as compatible on the page, contact IBM technical support representative for additional instructions.

9. Click **Finish** to initiate download of the drive firmware to each compatible drive that you selected in step 8.
10. When the Download Drive Firmware - Are you sure you wish to continue? window opens, type yes and click **OK** to start the drive firmware download. The Download Progress window opens. Wait until the download process is completed. Each drive that is scheduled for firmware download is designated as in progress until successful or failed.
11. If a drive is designated as failed, complete the following steps:
 - a. Click **Save as** to save the error log.
 - b. On the Subsystem Management window menu bar, click the menu option to display the storage subsystem event log and complete the following tasks that are necessary to save the event log before you contact your IBM service representative and proceed to the next step.
 - 1) Click **Select all**.
 - 2) Click **Save the Storage Subsystem Event Log**.
12. When the **Close** button becomes active, the drive firmware download process is complete. Click **Close** to exit the Download Progress window.
13. Use either of the following procedures to determine or verify which level of drive firmware is on a particular drive:
 - Right-click the drive on the Logical or Physical page in the Subsystem Management window and click **Properties**. The associated drive firmware version is listed in the drive properties table.
 - Select **Monitor > Reports > Storage Subsystem Profile** on the Logical or Physical page of the Subsystem Management window.

Storage Manager premium features

Storage Manager supports the following premium features, which are available separately for purchase from IBM or an IBM Business Partner:

Copy Services

The following copy services are available with Storage Manager:

- Enhanced FlashCopy and FlashCopy
- VolumeCopy
- Enhanced Remote Mirror Option
- Enhanced Global Mirror Option

For more information about the Copy Services features, see the *IBM System Storage DS Storage Manager Copy Services User's Guide*.

Storage Partitioning

Storage Partitioning is standard on all storage subsystems that are

supported by DS3000, DS4000, and DS5000 controller firmware versions. For more information about Storage Partitioning, see the “Storage partitioning overview” on page 55.

FC/SATA Intermix premium feature

The IBM System Storage DS3000, DS4000, and DS5000 Fibre Channel and Serial ATA Intermix premium feature supports the concurrent attachment of Fibre Channel and SATA storage enclosures to a single storage subsystem controller configuration. You can use this Intermix premium feature to create and manage distinct arrays that are built from either Fibre Channel disks or SATA disks and allocate logical drives to the appropriate applications, using a single storage subsystem.

For important information about using the Intermix premium feature, including configuration, firmware versions that are required for specific Intermix configurations, and setup requirements, see the *IBM System Storage DS Storage Manager Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview*. See your IBM representative or reseller for information about future storage subsystem support for the FC/SATA Intermix premium feature.

Hardware enhancement premium features

These are the premium features that provide support for certain storage hardware configurations, such as the following examples:

- a certain number of drive slots supported in the storage subsystems
- certain types of storage enclosures (for example, EXP810 attachment, or all eight EXP5060 attachments)
- certain types of drives (for example, solid state drives)

This premium feature category also includes the controller performance enhancement premium feature.

Full Disk Encryption

The use of full disk encryption (FDE) secures data against threats when an FDE drive is out of its owner's control.

Complete the following tasks to enable a premium feature on your storage subsystem:

- “Obtaining the premium feature enable identifier” on page 52
- “Generating the feature key file” on page 53
- “Enabling the premium feature” on page 53

Note: The procedure for enabling a premium feature depends on your version of the Storage Manager.

- “Disabling premium features” on page 54

To obtain the storage subsystem premium feature identifier string, make sure that your controller unit and storage enclosures are connected, the power is turned on, and they are managed using the SMclient.

Enabling the premium feature trial version

Note: This trial version is offered only for the DS3500 and DCS3700 storage subsystems, and the DCS3700 storage subsystems with Performance Module Controllers.

IBM Storage subsystems with controller firmware version 7.83 and later provide trial versions of certain premium features for up to 90 days. Alerts are sent after completion of 30 days, and 3 days before the trial period expiry. The permanent premium feature can be purchased and activated any time during the 90 days trial. If the permanent premium feature is activated during the trial period, the storage subsystem configurations created with the premium feature trial continue to be valid. If you do not purchase the premium feature, you must delete the storage subsystem configurations you created using the premium feature trial version. Else, the storage subsystem will be placed in 'out-of-compliance' state for the trial premium feature. To enable premium feature trial version, click **Try Now** in the **Premium Feature and Feature Pack** window.

Enabling the permanent premium feature

You must obtain the permanent premium feature, generate a unique key for it, and then enable the permanent Premium feature. You can also disable the permanent premium feature.

Obtaining the premium feature enable identifier

Each storage subsystem has its own unique premium feature enable identifier. This identifier ensures that a particular feature key file is applicable only to that storage subsystem.

Before you obtain the feature enable identifier, complete the following prerequisites:

1. Make sure that you have available the feature activation code from the premium feature web activation card, as well as the model, machine type, and serial number of the storage subsystem.
2. Make sure that the controller unit and storage expansion enclosures are connected, turned on, and configured.

To obtain the feature enable identifier, complete the following steps:

1. Click **Start > Programs > Storage Manager xx Client**. The Enterprise Management window opens.
2. In the Enterprise Management window, double-click the storage subsystem for which you want to enable the premium feature. The Subsystem Management window opens for the selected storage subsystem.
3. Complete one of the following actions, depending on your version of the Storage Manager:
 - If you are using Storage Manager version 9.x or earlier, click **Storage Subsystem > Premium Features > List**. The List Premium Features window opens and displays the feature enable identifier.
 - If you are using Storage Manager version 10.x or later, click **Storage Subsystem > Premium Features....** The Premium Features and Feature Pack Information window opens. The feature enable identifier is displayed at the bottom of the new window.
4. Record the feature enable identifier.

Note: To prevent a mistake when recording the feature enable identifier, copy the 32-character identifier and paste it in the premium feature key request field.

5. Click **Close** to close the window.
6. Continue to “Generating the feature key file” on page 53.

Note: To check the status of an existing premium feature in Storage Manager version 9.x or earlier, select **Storage Subsystem > Premium Features > List** from the menu.

Generating the feature key file

You can generate the feature key file with the Premium Feature Activation tool that at <http://www.ibm.com/storage/fasttkeys>.

1. Complete the steps on the website.

Note: Make sure that you select the correct Premium Feature or Feature Pack after you are prompted.

The Feature key file is available for download on the webpage, and it can be emailed to you.

2. On your hard disk drive, create a new directory (for example, name the directory FlashCopyfeaturekey).
3. Save the premium feature key file in the new directory.

If the premium feature key is lost, or if the premium feature identifier is changed and the premium feature is no longer in compliance, you can request a premium feature reactivation key file at <http://www.ibm.com/storage/fasttkeys>. You must have the same machine type, model, and serial number information available that you used to generate the premium feature key files initially.

Enabling the premium feature

To enable the premium feature, follow the applicable procedure for your version of Storage Manager.

Enabling the premium feature and feature pack in Storage Manager 10.x or later

To enable a premium feature in Storage Manager version 10.x or later, complete the following steps:

1. In the Subsystem Management window, click **Storage Subsystem > Premium Features....** The Premium Features and Feature Pack Information window opens.
2. To enable a premium feature from the list, click **Enable** or **Use key file**, depending on the version of the controller firmware. A window opens that allows you to select the premium feature key file to enable the premium feature. Follow the on-screen instructions.
3. Verify that the premium feature is enabled by inspecting the displayed list of premium features in the Premium Features and Feature Pack Information window.
4. Click **Close** to close the Premium Features and Feature Pack Information window.

Enabling the feature pack

1. Click **Change** in the Premium Feature and Feature Pack Information window.
2. A window opens for you to select the feature pack key file. Select the key file and click **OK**.
3. Review the contents of the **Feature Pack installed on storage subsystem** field to verify whether the feature pack is installed.

Important: Enabling a premium feature pack requires that the controllers be restarted. If the storage subsystem for which the premium feature pack will be enabled is running, be sure to schedule downtime to restart the controllers.

Disabling premium features

In normal system-operating conditions, you do not have to disable the premium features. However, if you want to disable a premium feature, make sure that you have the key file or the premium feature entitlement card with the premium feature activation code for generating the key file. You will need this key file to re-enable the premium feature at a later time.

Note:

1. If you want to enable the premium feature in the future, you must reapply the Feature Key file for that feature.
2. You can disable the Remote Mirror Option without deactivating the feature. If the feature is disabled but activated, you can perform all mirroring operations on existing remote mirrors. However, when the feature is disabled, you cannot create any new remote mirrors. For more information about activating the Remote Mirror Option, see the *IBM System Storage DS Storage Manager Copy Services User's Guide* or see "Using the Activate Remote Mirroring Wizard" in the Storage Manager online help.
3. If a premium feature becomes disabled, you can access the website and repeat this process following the re-activating premium features option.

Disabling the premium feature in Storage Manager 10.x or later

To disable a premium feature in Storage Manager version 10.x or later, complete the following steps:

1. In the Subsystem Management window, click **Storage Subsystem > Premium Features**. The Premium Features and Feature Pack Information window opens.
2. Select the premium feature that you want to disable and click **Disable**.

For additional assistance, contact your local IBM service provider.

Saving the storage subsystem profile

Important: You must save a storage subsystem profile whenever you modify the arrays and logical drives in your storage subsystem. This saved profile contains detailed controller information, including logical and physical disk configuration information, that can help you recover the configuration in the event of a catastrophic failure. Do not save a profile for a storage subsystem on that same storage subsystem.

To save a storage subsystem profile, select the menu option to display the subsystem profile in the System Management window and click **Save As** when the Storage Subsystem Profile window opens. To save the full profile, select the **All** tab. You can also select the menu option to save the support data to collect all the various types of inventory, status, diagnostic and performance data from this storage subsystem and save them in a single compressed file.

Chapter 4. Configuring storage

After Storage Manager is installed, you must configure the storage subsystem or subsystems. The following topics in this chapter describe the tasks that are necessary for configuration:

- “Storage partitioning overview”
- “Using the Task Assistant” on page 56
- “Configuring global hot-spare drives” on page 80
- “Configuring disk storage” on page 67
- “Defining a default host operating system” on page 81
- “Defining a host group” on page 83
- “Defining heterogeneous hosts” on page 83
- “Defining the host and host ports” on page 84
- “Mapping LUNs” on page 84
- “Using Performance Read Cache” on page 88

Near the end of this chapter, the following topics provide optional information that might apply to configuring your storage subsystems:

- “Configuring the IBM System Storage for IBM i” on page 85

Note: This section applies only to storage configurations that use the IBM i operating system.

- “Configuring and using optional premium features” on page 86

Note: This section applies only to storage subsystems that have premium features.

- “Using other features” on page 89
- “Tuning storage subsystems” on page 95

Note: By default, the **Setup** tab in the Enterprise Management window opens first when you start Storage Manager. See “Enterprise Management window” on page 13 for a detailed description of the Enterprise Management window.

Storage partitioning overview

Before you create storage partitions, be aware of the following information:

- The Storage Manager Task Assistant provides a Storage Partitioning wizard that you can use to define your host and host ports and map LUNs to the storage partitions. If your storage subsystem is running controller firmware 05.xx.xx.xx, you cannot use the wizard. Both types of procedures are documented in this section.
- These procedures assume that you have already created a physical connection between the host and the storage subsystem controllers, and that you have also connected and zoned the switch (if applicable). If you have not completed these connections, Storage Manager cannot list the worldwide port names (WWPNs) or the iSCSI iqn-names of the HBAs during these procedures. In this case, you must type the WWPNs into the applicable fields during the procedure that is described in “Defining the host and host ports” on page 84.

- Create the host group at the storage subsystem level. Do not create host groups at the default group level.

Note: If you have a DS4100 or a DS4300 configuration and partitioning is not enabled, you can use the default host group.

- In case of multiple HBAs in the host connected to the storage subsystem, create a single host partition to include all. Use the host group definition only to group a set of hosts that share the same set of logical drives.
- In a cluster partition, perform logical drive mappings on the host group level so that all of the hosts can recognize the same storage. In a normal partition, perform logical drive mappings on the host level.
- To set up and assign IBM i LUNs on the DS5300 and DS5100 storage subsystems with the wizard, see “Configuring the IBM System Storage for IBM i” on page 85 for information specific to IBM i configuration.

Using the Task Assistant

The Storage Manager Task Assistant provides a central location from which you can choose to perform the most common tasks in the Enterprise Management window and in the Subsystem Management window. You can use the Task Assistant to complete many of the procedures that are described in this section.

Important: If you have controller firmware version 7.50 or later, the Storage Manager task descriptions might differ slightly from the tasks in the following lists.

In the Subsystem Management window, the Task Assistant on the Setup tab (depending on the firmware controller version) consists of shortcuts to these tasks:

- Configuring storage subsystems
- Defining hosts
- Creating a new storage partition
- Mapping additional logical drives
- Saving configurations

If there is a problem with the storage subsystem, a shortcut to the Recovery Guru is displayed. Use the Recovery Guru to learn more about the problem and find solutions to correct the problem.

Important: If you have controller firmware version 7.50 or later, the Storage Manager procedure for accessing the Task Assistant functionality is slightly different. There is no button and no separate window for Task Assistant. Click the **Setup** tab in the Subsystem Management window to access the Task Assistant menu in the Initial Setup Tasks window.

Note: The Task Assistant is automatically invoked every time you open the Subsystem Management window, unless you select the **Don't show the task assistant at startup again** check box at the bottom of the window.

Drives supported by IBM System Storage DS Storage Manager

The DS subsystem supports hard disk and solid state disk drive media types.

Depending on the models, the DS subsystem supports some or all of the following disk drive interfaces:

SATA In certain enclosures, the SATA disk drive requires an ATA translator or FC-SATA interposer that helps you insert the drive into the drive slot with a FC connector.

Note: The SATA drives and the ATA translator or FC-SATA interposers are sold as a single identity.

Fibre-Channel (FC)

There are no special requirements for FC drives.

SAS For enclosures with FC mid plane, the SAS drive requires the FC-SAS interposer that helps you insert into a drive slot with a FC connector. This category also includes NL SAS drives.

Note: The SAS drives and the FC-SAS interposers are sold as a single identity and is known as FC-SAS drive.

In addition to the differences in the disk drive types and interfaces, there are few differences with respect to the drive capabilities like T10 Protection Information (T10PI) or Full Disk Encryption/Self-Encryption (FDE/SED) capabilities. Drive capacities are available for most of the supported drive media types, drive interfaces, and drive capabilities. The DS subsystem does not support all types of drive media. For more information on the type of drives available and supported for a given storage subsystem, refer to **DS subsystem RFAs**. You can also refer the DS storage subsystem Installation, User's and maintenance guide for more information on the FRU part list for the drives supported in a storage subsystem model. A summary of supported drive type, drive interface and drive capability is shown in Table 11.

Table 11. Summary of supported drive types, interfaces and capabilities

Drives supported by the IBM DS Storage Manager	Drive Interface	T10 PI capability	SED capable	Non-SED capable
Hard Disk Drive	SATA	N/A	N/A	N/A
	Fibre-Channel (FC)	Yes	Yes	Yes
		N/A	Yes	Yes
	NL SAS/SAS	Yes	Yes	Yes
		N/A	Yes	Yes
Solid State Drive	FC	N/A	N/A	Yes
	SAS	N/A	N/A	Yes

In the Subsystem Management, the Physical or Hardware tab has buttons or drop-down list that helps identify various drive types in a given enclosure depending on the controller firmware version. When you click the button or select from the drop-down list, all drives that meet the button definition are highlighted in the physical view pane. In the enclosure, if all the drives meet the button definition, the button will be disabled.

Rules for selecting drives when creating a RAID array

These are the things to consider when selecting drives for a RAID array

- The RAID array can be created using only drives using the same drive interface.

For example, you cannot create an RAID array with FC, SAS and SATA drives even though the SATA and SAS drives have interposers to allow them to behave like FC drives. Also, a drive with a specific drive interface cannot be used as a spare for drive in a RAID array with a different drive interface.

- A RAID array cannot have a combination of spinning hard disk drives and solid state drives.
- Because NL-SAS and SAS drives operate at a different rotational speeds, do not mix the NL-SAS and SAS drives in the same RAID array. The NL-SAS operates at 7200 rpm and the SAS drives operate at 10K or 15K rpm. In a RAID array composed of SAS drives and NL-SAS drives, it is possible to use a NL-SAS to replace a failed SAS drive or SAS drives to replace a failed NL-SAS drive.
- Do not mix drives of the same drive interface that operate at different rotational speeds in the same RAID array. The Storage Manager GUI or SM command line interface does not prevent this configuration. If a suitable drive is not available, the controller firmware might select an available hot-spare drive with a different rotation speed as a spare for the failed drive. In case the hot spare drive has a lower rotation speed, replace the failed drive as soon as possible.
- If a drive of required rotational speed is not available, IBM might supply a similar drive with higher rotational speed as a replacement FRU. The performance of the array or a disk pool is not affected when a higher rotational speed drive is used as a replacement.
- In a RAID array, drives of different capabilities like T10PI or FDE can mix with the drives without these capabilities in the same RAID array only if one of the drive capabilities is not enabled for that RAID array. For example, you can create a RAID array with both T10PI supported drives and non-T10PI supported drives. However, the created array will not be able to operate with T10PI functionality enabled.
- Drives with lower rotational speed can be used as spares for array with drives having higher rotational speeds. It is recommended that you do not mix drives with different rotational speeds in the same RAID array because the performance of an array or a disk pool might be degraded with lower rotation speed drives.
- Drives with different sizes can be mixed in a RAID array. However, the array will be created with all drives having size of the smallest drive size in the RAID array.
- A RAID array with additional capabilities enabled, like FDE and T10PI, cannot have drives without the enabled capabilities used as spares for a failed drive in the RAID array. For example, a RAID array with T10PI and FDE enabled requires a drive that has T10PI and FDE capabilities as a hot-spare drive.
- Disk pools can be composed from spinning hard disk drives with a SAS disk interface only. Disk pools support Enhanced FlashCopy, Metro Mirroring, and Global Mirroring, and do not support Legacy FlashCopy. A single disk pool cannot have drives of different capacities and different spindles.
- 2 TB NL-SAS drives (FC#3450, 3451) do not support T10 PI. If you are using these drives on a DCS3700 subsystem with Performance Module Controllers, you must manually create an array or a disk pool to avoid T10 PI enabling. To manually create a disk pool, click **No** in the **Disk Pool Automatic Configuration** window and clear **Filter drive selection to show T10 PI (T10 Protection Information) capable drives only**. To create a T10 PI-capable logical drive in an array, all drives must be T10 PI capable.

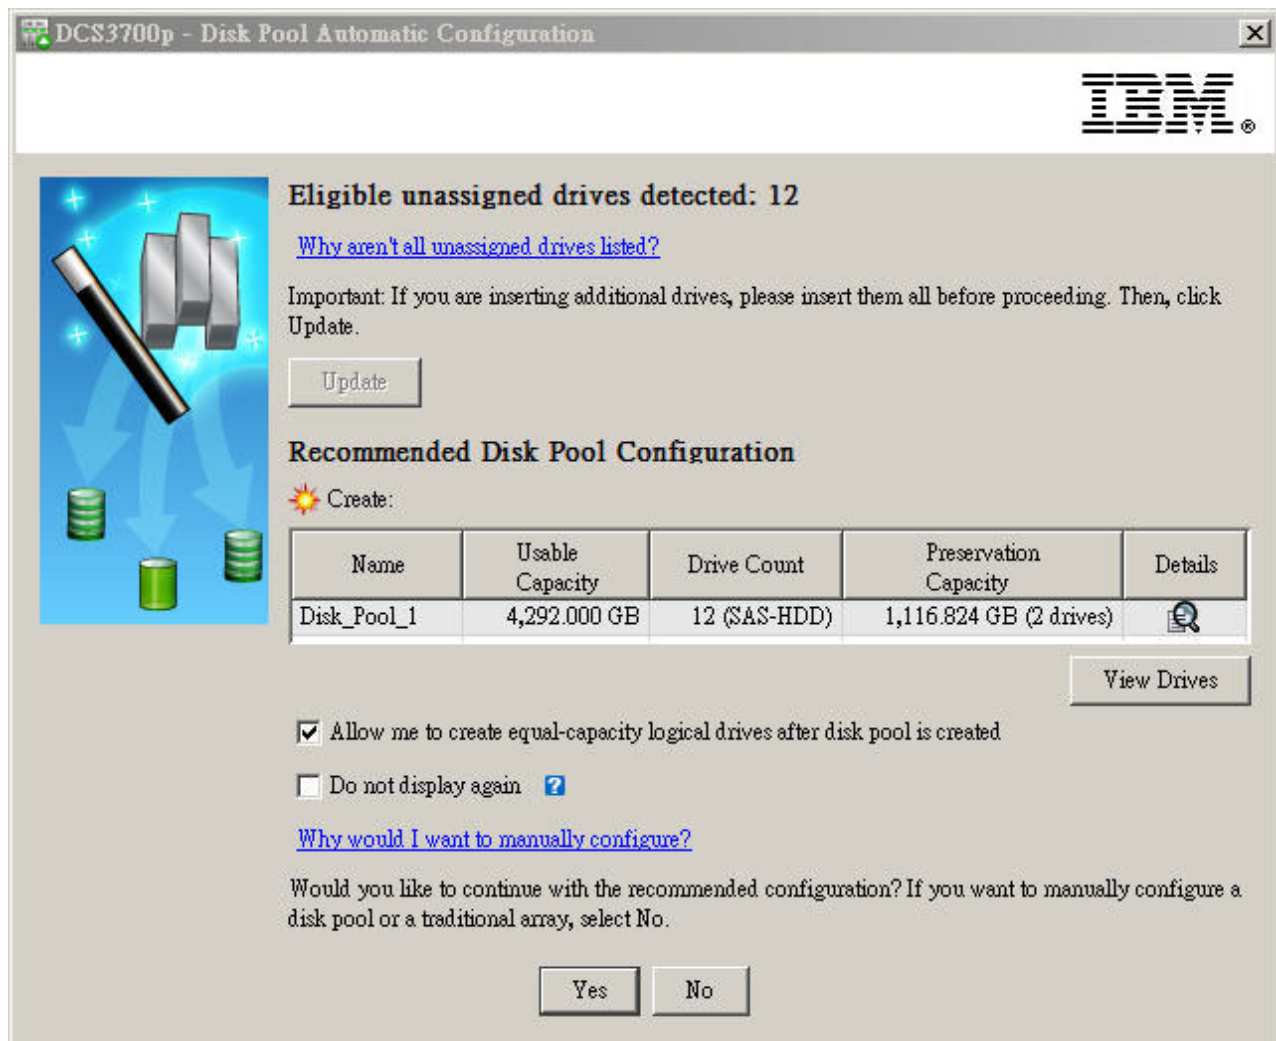


Figure 3. Disk Pool Automatic Configuration

Solid state drive (SSD) attributes

Solid State drives are drives that store data in a flash memory chip instead of on rotating hard disk surfaces resulting in faster random access write/read speeds than hard disk drives.

Because flash devices have certain write cycle count limitations, these drives contain additional spare flash memory chips to support write cycle use during the drive warranty period. shows an SSD disk drive and the percentage of the remaining spare blocks in the SSD drive. SSD is supported with controller firmware version 7.60.xx.xx and later because it supports media scan for logical drives in RAID arrays composed of SSD drives. However, controller firmware version 7.77.xx.xx or later is recommended.

For subsystems with controller firmware version 7.8x.xx.xx or later, the DS5100/5300 and DS5020 storage subsystems can have a maximum of 20 SSD drives in a given storage subsystem configuration. The DS3500 and DCS3700 storage subsystems, and DCS3700 storage subsystems with Performance Module Controllers can have a maximum of 24 drives. Additionally, for the DS5100/5300

subsystems, you should not have more than four SSDs per drive expansion enclosure and you should spread the SSDs among the drive enclosures on as many different drive channels as possible. This is to ensure optimal performance of your storage subsystem.

T10PI capable drive attributes

T10PI capable drives support the industry standard extension T10 Protection information (T10 PI or T10PI) specified for the SCSI protocol. T10PI standard is also referred to as T10 DIF (Data Integrity Field).

This standard is designed to provide an extra level of data integrity by safeguarding each user data block with 8 bytes of integrity metadata while the data is transferred between the storage controller and the T10PI-initialized disk drives. This extra level of data integrity is also extended to the server when the storage subsystem is I/O attached to IBM Power servers running the AIX operating system. T10 PI support within AIX operating system includes the Protection Information metadata with user data transferred between the AIX server and the storage controller. T10PI standard is designed to provide the user with end-to-end protection and correction against silent data corruption due to device driver errors, filesystem errors or misdirected, lost, phantom, or torn writes.

The T10PI capable drives are initialized as Type 2 T10PI drives. These drives have 520 byte sectors instead of the standard 512 byte sectors. The additional 8 bytes contains the Protection Information metadata described in Table 12, that can be used to verify the data in-flight and at rest.

Table 12. Protection Information metadata (8 bytes)

Byte #1	Byte #2	Byte #3	Byte #4	Byte #5	Byte #6	Byte #7	Byte #8
Logical Block Guard (2 bytes) <ul style="list-style-type: none"> • 16-bit CRC • Receiver computes CRC on received data and compares to received CRC • Protect the data portion of the sector 		Logical Block Application Tag (2 bytes) <ul style="list-style-type: none"> • May be owned by application client (initiator) or device server (target) • Checked only by the owner 		Logical Block Reference Tag (4 bytes) <ul style="list-style-type: none"> • Receiver optionally checks against expected value • Protect against out-of-order and misdirected write scenarios 			

The DS storage subsystem supports the T10PI Type1 host protection scheme. Figure 4 on page 61 shows where Protection Information metadata is checked from the application in the host to the drive in the storage subsystem.

Protection Information (PI) Check Points

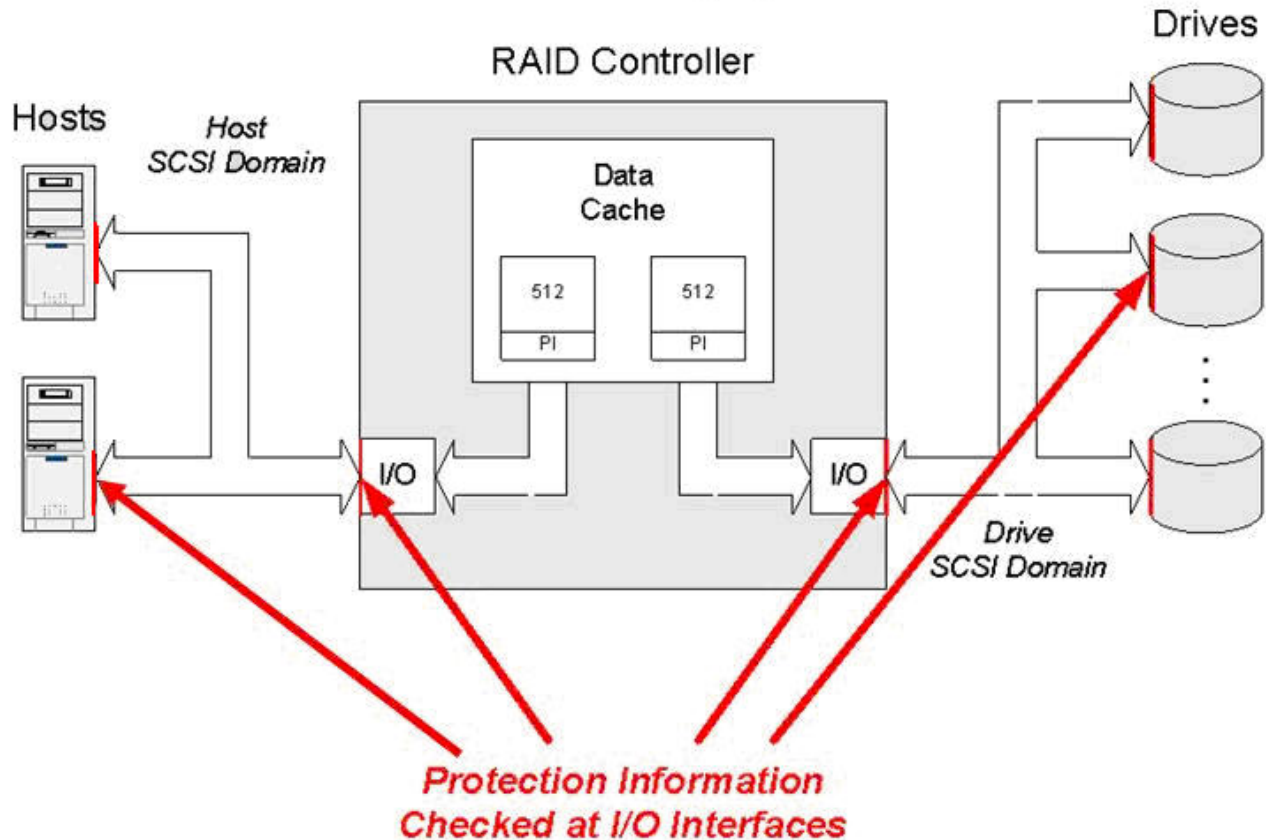


Figure 4. Protection Information (P) check points

Rules for enabling T10PI functions

To enable T10 PI functions in a DS storage subsystem, ensure that the following conditions are met.

T10PI functionality supported controller firmware

The T10PI functionality is enabled by the controller firmware instead of a premium feature. Controller firmware that supports T10PI is 7.77.xx.xx or later. Currently, DS3950, DS5020, DCS3700, and DS5100/DS5300 subsystems only support T10 PI. A DCS3700 storage subsystem with Performance Module Controllers supports T10 PI if the controller firmware is 7.83.xx.xx or later. A DS3500 storage subsystem supports T10PI if the controller firmware is 7.86.xx.xx. Contact IBM resellers or representatives for support of T10PI on other DS storage subsystems in the future.

Note: T10PI drives can be used in the storage subsystem with controller firmware that does not support T10PI functionality. In such cases, the T10PI drives are treated as non-T10PI capable drives.

Fibre-Channel Host Interface

Fibre-Channel Host Interface must be installed on the DS storage subsystem. Additionally, the T10PI-enabled logical drive must be mapped to a host port that is discovered through the DS storage subsystem Fibre-Channel port. For example, if the controller has FC, iSCSI, or SAS

host interface installed at the same time, the T10PI-enabled logical drives can be mapped to the host port that is discovered through the FC port only. An error occurs if you attempt to map a T10PI-enabled logical drive to a host port that is discovered through the iSCSI or SAS interface.

(For an AIX server) NVSRAM file with T10PI-enabled host type

The NVRAM files provided with the controller firmware version 7.77.xx.xx or later must be installed in the storage subsystem. These NVSRAM files have T10PI enable bit set in the AIX and AIXAVT host type regions to extend T10PI functionality to the server. Without this bit set in the host type region, T10PI functionality can be enabled between the subsystem controller and drives only.

Note: Refer to the SSIC for information about the types of FC adapters supported along with the required device driver, firmware version, and the versions of AIX operating systems that provide T10PI support at the server.

Creating a T10 PI capable array

Complete the following steps before creating an array which is T10 PI capable.

1. Click **Total Unconfigured Capacity**. The **Create Array** window opens.
2. If you are using multiple types of drives, select **HDD-SAS** in **Drive type**.
3. Select the check box **Filter drive selection to show T10 PI (Protection Information) capable drives only**.
4. Go to “Creating an array” on page 74.

Creating a T10 PI capable disk pool

Complete the following steps before creating a disk pool which is T10 PI capable.

1. Click **Total Unconfigured Capacity**.
2. If you are using multiple types of drives, select **HDD-SAS** in **Drive type**.
3. Select **Only T10 PI-capable drives** in **Data assurance**.
4. Go to “Creating a disk pool” on page 72.

Enabling and disabling a T10PI capable RAID array

To create a T10PI-capable RAID array, all of the drives in the RAID array must be T10PI capable.

This task assumes that you are using the subsystem management interface to define the array.

To enable T10PI capability on a RAID array:

Note: The screenshots in this section are only illustrative and may differ from the actual UI depending on the Storage Manager and controller firmware versions.

- In the Specify Capacity/Name (Create Logical Drive) dialog click **Enable T10 PI (Protection Information) (T10 PI) on the new logical drive** as shown in Figure 5 on page 63.

Note: If you create the logical drive without T10PI functionality, it cannot be converted into a T10PI-enabled logical drive later.

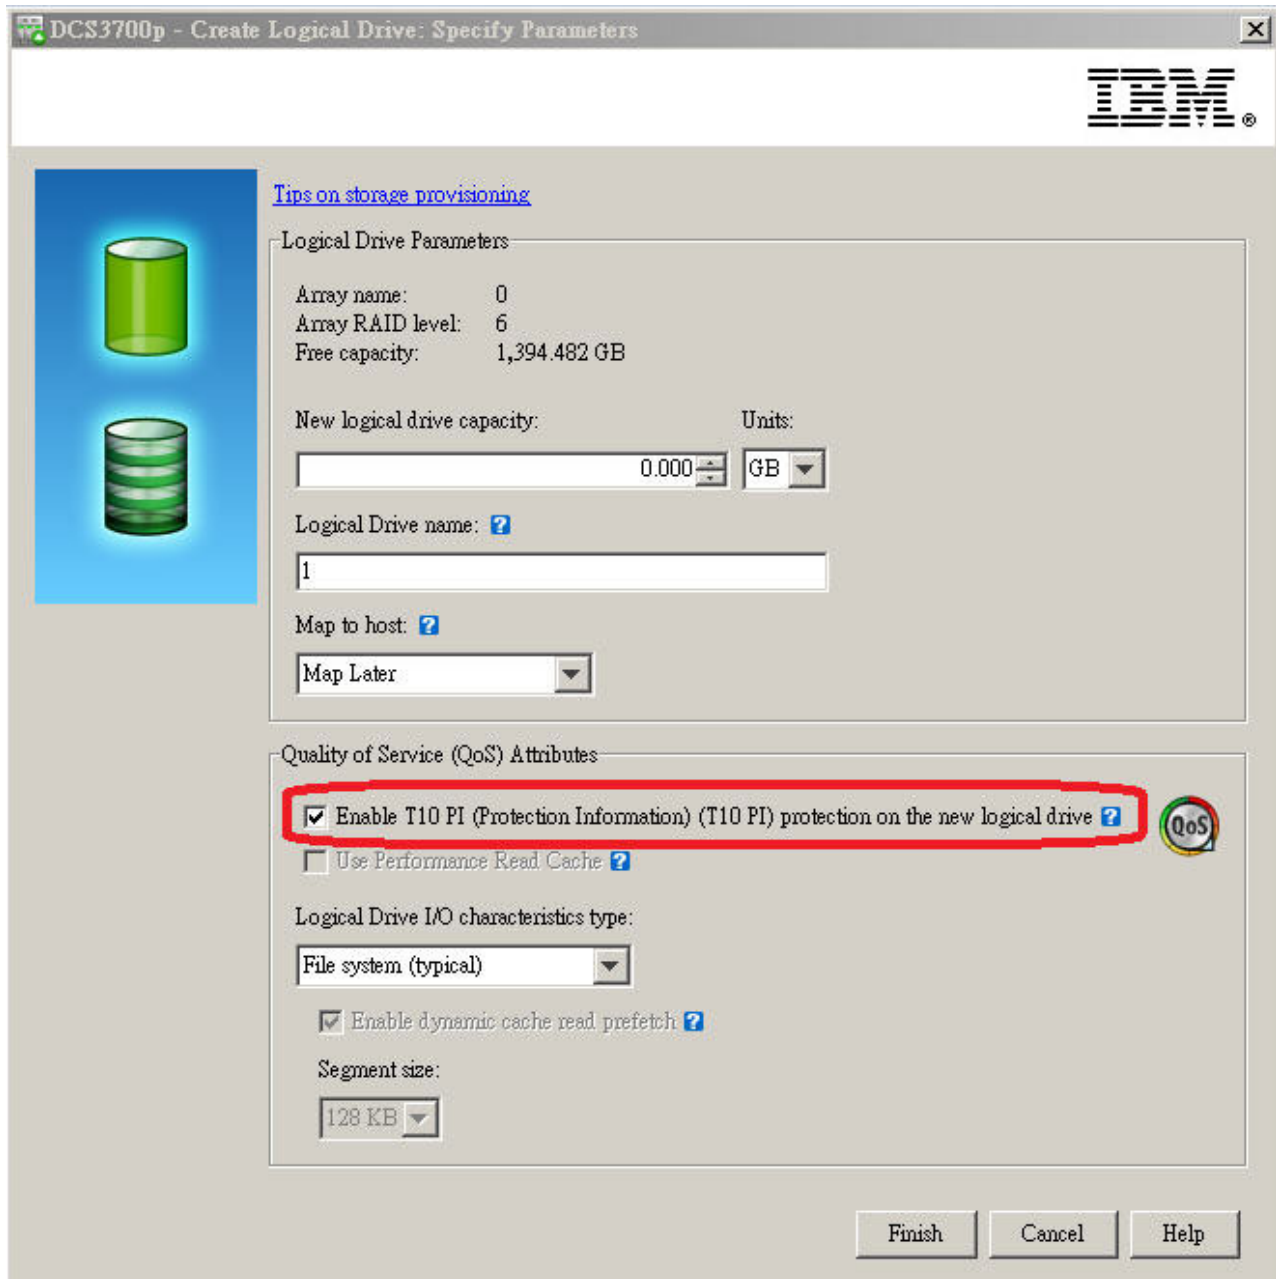


Figure 5. Enabling T10 PI on a logical drive

Note: An additional parameter also exists on the appropriate SMcli commands to indicate whether a logical drive is created with T10PI enabled. Figure 6 on page 64 shows a RAID array and its logical drive that has T10PI functionality enabled. The shield icon indicates that the array is a T10PI capable RAID array.

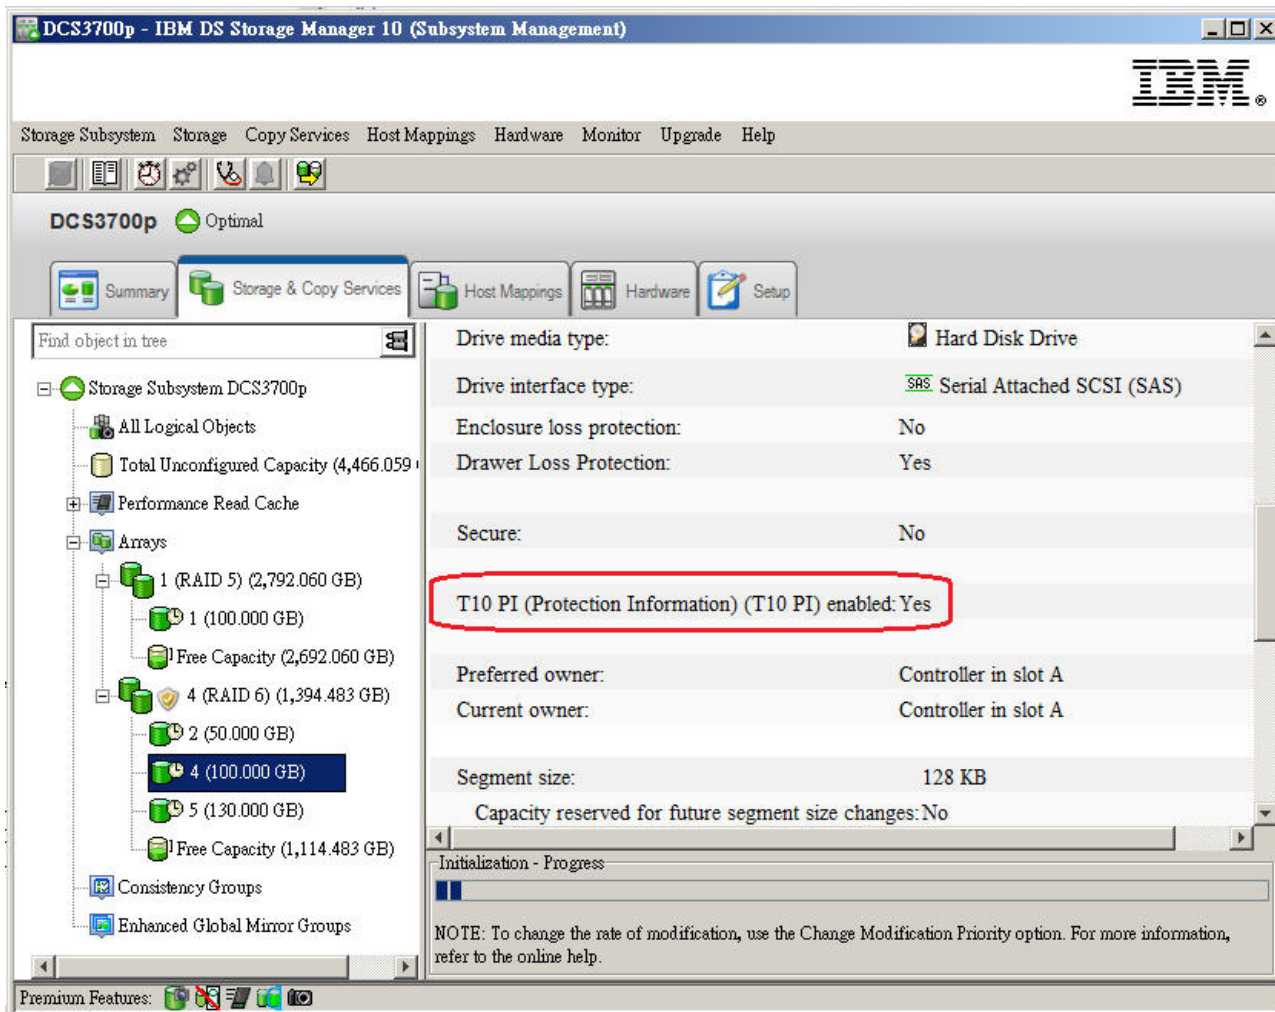


Figure 6. RAID drive - Protection Information (T10 PI) - enabled

Note: You do not have to create all logical drives in a T10PI-capable RAID array with T10PI enabled. For example, logical drive 4 of RAID array 4 is not T10PI enabled but logical drives 2 and 5 are T10PI enabled as shown in Figure 7 on page 65. However, because you can only enable T10PI functionality at the time of creation, it is recommended to create a logical drive with T10PI enabled and then disable it at later time, if necessary.




Figure 7. Example - Logical Drive 4 of RAID array 4 - T10PI not enabled

To disable T10PI capability for a RAID array:

- Right-click the drive that you are disabling T10PI capability and select **Disable T10PI (Protection Information) (T10PI)** as shown in Figure 8 on page 66.

Notes:

- There is also an SMcli command to disable T10PI capability.
- Once you disable T10PI capability on a logical drive, you cannot simply enable T10 PI on the same logical drive; you must delete the drive and re-create it with T10PI enabled.




Figure 8. Disabling T10PI

Full Disk Encryption (FDE) attributes

The Full Disk Encryption (FDE) or Self-encryption (SED) drives have built-in encryption mechanisms to protect the drive information from unauthorized access from outside the DS storage subsystem.

The secured FDE/SED drive is locked when power is switched on. Unlock the drive by sending the appropriate security key from the controller to the drive. Manage the security key locally inside the controller or through an external key

manager like IBM Tivoli Key Lifecycle Manager (TKLM). For additional security, the drive encrypts before it is written preventing the data on the disk surface from being scanned when the drive fails or is removed from active use. The FDE capable RAID array/disk pool is shown with either the unlock icon if the array/disk pool is not secured and with the lock icon if the array/disk pool is secured; as shown in Figure_10. For more information about FDE drives, see Chapter 6, "Working with full disk encryption," on page 191.




Figure 9. FDE capable RAID array - security details

Configuring disk storage

You can use the IBM DS Storage Manager version 10.83 with controller firmware version 7.83 to configure a collection of drives into either a disk pool, an array, or both. A disk pool or an array contains drives with the same or similar characteristics. The characteristics used to determine similar drives are:

- **Drive type** — The types of drives supported for an array are Fibre Channel, SATA, FC-SAS, or SAS. A disk pool can consist of only SAS drives.
- **Drive media type** — The drive media supported for an array is Hard Disk Drive (HDD) or Solid State Disk (SSD). A disk pool can only have SAS HDD drives. SAS SSD drives are not supported in a disk pool.

Table 13. Types of drives that can be used in arrays and disk pools

Drive Types	Array	Disk Pool
SAS Disk	Yes	Yes
SAS Disk - T10 PI	Yes	Yes
SAS Disk - SED	Yes	Yes
SATA Disk	Yes	Yes
FC Disk	Yes	No
FC Disk - SED	Yes	No
FC-SAS Disk	Yes	No
SSD	Yes	No

- **Rotation speed** — The rotation speed of the drives in a disk pool or an array must be the same. It is possible to create arrays or disk pool with drives of different rotation speeds, but such a configuration will not be optimal.

Note: You cannot create a disk pool with drives of different rotation speed from the Storage Manager User Interface. You will have to use a SMcli command to create a disk pool comprising drives having different rotation speed.

- **Security level** — The drive security must be the same for all the drives if an entire disk pool or an array is to be secured at a designated security level.
- **Capacity** — To efficiently use the drives in a disk pool or an array, the capacity of these drives should be the same. If the drives in a disk pool or an array have different capacities, the storage management software will only use as much capacity as that of the smallest drive in the disk pool or array. For example, if the disk pool comprises several 4 GB and 8 GB drives, the DS Storage Manager will use only up to 4 GB on each drive, leaving 4 GB of the 8 GB drives unused.

An *array* is a set of Fibre Channel, SAS, SATA, or solid state drives that are logically grouped together to form a Redundant Array of Independent Disks (RAID). Depending on the needs, different arrays can be created in a storage system with different RAID characteristics: RAID levels, segment size, and stripe width. Arrays can be either standard or secured (with full disk encryption).

A *Disk Pool* is an alternative method of organizing the SAS disks in the storage subsystem into logical disk units and presenting them to the host. Disk pools can be either standard or secured (with full disk encryption). You can use the IBM DS Storage Manager version 10.83 with controller firmware version 7.83 to configure a collection of drives into either a disk pool, an array, or both. Currently, DS3500 and DCS3700 storage subsystems, and the DCS3700 storage subsystems with Performance Module Controllers support disk pooling, if the controller firmware version is 7.83 or later.

Table 14. Copy Services support by array and disk pool

Copy Service	Array	Disk pool
FlashCopy Logical Drive	Yes	No
Enhanced FlashCopy	Yes	Yes
VolumeCopy	Yes	Yes
Remote Mirroring	Yes	Yes
Global Mirroring	Yes	Yes

Comparisons between disk pool and array

A disk pool is different from the array in many aspects.

RAID protection - An array can be created using one of these RAID levels: 0, 1, 10, 3, 5, or 6. The data and parity segments are striped over the drives as defined by the RAID level. Table 14 contains information about the RAID level that should be selected for an array. In the array, the RAID stripe covers a segment in each of the drive that is part of the RAID array. There is almost a direct relationship between the LBAs of the RAID arrays and the physical location of data on the disks in the array.

You cannot specify the RAID level for a disk pool. Disk pools are divided into 4 GB chunks, which are called D-chunks, of disk space. Each D-chunk consists of ten 512 MB pieces, which are called D-pieces. Ten D-pieces of a D-chunk must be written on 10 different disks. If that is not possible, the disk pool will be degraded (if any of the ten pieces cannot be written on a separate disk) or failed (if two or more of the ten pieces cannot be written on a separate disk). Each D-piece consists of 4096 segments of 128 KB. A segment from each of the ten D-pieces forms a RAID 6 8D+P+Q stripe. In different words, within the 4 GB D-chunk, the data and RAID parity are stored in 4096 of 8D+P+Q RAID-6 stripes. The data is stored contiguously with the RAID stripe and the D-chunk. The logical drives in the disk pool are created with one or more 4 GB chunks of disk space. If there are more than 11 disks in a disk pool, the data will be distributed to all the drives as multiples of 4 GB chunks. Because of the pattern of allocation of D-chunks among disk drives, a clear, direct relation does not exist between the LBAs of the disk pool and the physical locations of data on the disks. The more the drives on a disk pool, the more the data distributed among the drives.

RAID array characteristics - When you create an array, you can dictate what a RAID stripe should be, depending on the number of drives in an array, the RAID level, and the segment size of the data in each drive. In some applications, an array can be created with a RAID stripe by tuning the segment size and the RAID stripe width. This enables full-stripe write feature of the controller that might result in better write performance. However, you cannot tune the segment size and the RAID strip width characteristics of a disk pool in the current implementation. Disk pool can be created with more than the minimum prescribed 11 drives. However, the RAID data that are distributed among the drives is always having the RAID characteristic of RAID 6 8D+P+Q with segment size of 128 KB.

Hot-sparing - In the event of a drive failure, arrays use dedicated, global hot-spares for array reconstruction, whereas in disk pools, a certain percentage of each drive in the disk pool is reserved for disk pool reconstruction. Table 15 on page 70 shows the drive count equivalent of reserved drive capacity for various disk pool sizes. The advantages reserved capacity for reconstruction or hot-sparing include:

1. *Improvement in reconstruction time* - Because each drive in a disk pool has reserved capacity for reconstructed data, all drives in a disk pool can be used to write reconstruction data. In the case of array and hot-spare drives, only one drive is available for writing reconstructed data, causing a bottleneck in the reconstruction process.
2. *No drive in standby*. There have been instances when hot-spare drives have failed during reconstruction because of grown disk-platter defects or other hardware defects, because of drive inactivity. By not dedicating any drives for reconstructed data, all the drives in the disk pool are used at any point in time.

3. *Improvement in copy-back time* When the failed drive is replaced, there is also some improvement in the copy-back time, because all the drives are engaged in copy-back. Reconstructed segments from various drives in the disk pool are retrieved and written to the replaced drives. In an array, the reconstructed data from one drive is copied back to the replaced drive.

Table 15. Reserved capacity in a disk pool

Number of drives in a disk pool	Number of drives reserved
11	1
12-31	2
32-63	3
64-127	4
128-191	6
192	7

Storage efficiency - If an array is divided into multiple logical drives, each logical drive is assigned a fixed slice of the array. Unused capacity in some logical drives cannot be used as additional capacity in other logical drives or to create new logical drives. The logical drives in a disk pool are managed together. There is no fixed allocation of capacity for each logical drive. This will allow a logical drive to use its share of the disk pool and offer the whole unused capacity of the disk pool for any logical drives in a disk pool that need it or to create additional logical drives. If an array is divided into multiple logical drives, each logical drive takes up a fixed portion of the array. The size of the logical drive can be increased as needed if there is available space in the array. However, the unused capacity in a logical drives cannot be freed up to be used to create new logical drives or as additional capacity in another logical drive in the array. The only way to free up the space in an existing logical drive is to back up the data in the logical drive, delete the logical drive, re-create the logical drive with a smaller capacity and then restore the backed up data into the smaller logical drive. If a disk pool is divided into logical drives, the logical drives can be created as thin logical drive, which has only a fraction of the logical drive space allocated physically on the disk pool hard disks at the time the logical drive is created. As the logical drive needs more physical space, additional space from the disk pool free capacity will be allocated to logical drive in 4 GB D-chunks, ensuring that the logical drive is correctly provisioned without any unused free capacity.

Table 16. List of features supported in an array or a disk pool.

Features	Array	Disk pool
Enclosure lost protection	Yes ¹	Yes ²
Drawer lost protection	Yes ¹	Yes ²
Dynamic Logical drive Expansion	Yes	Yes
Dynamic array/disk pool Expansion	Yes	Yes
Thin provisioning	No	Yes

Note:

1. The maximum number of drives of an array for enclosure or drawer lost protection depends on the RAID level of the array. For RAID 1, 10, 3 and 5

- arrays, it is one drive per enclosure or drawer and for RAID 6 array, it is two drives per enclosure or drawer for enclosure or drawer lost protection.
2. The maximum number of drives of a disk pool per enclosure or drawer for enclosure or drawer lost protection is 2.

Logical Drives

Within a disk pool or an array, the drives are further organized into logical drives. A logical drive is a logical component that a host uses to organize data storage on a storage subsystem. The host operating system sees a logical drive as a single drive even though data is written across several physical drives within a disk pool or an array. Each logical drive is mapped to a logical unit number (LUN) that a host uses to access a logical drive. A host attached to a storage subsystem writes data to the logical drives and reads data from the logical drives. Logical drives cannot span across multiple disk pools or arrays. A disk pool or array can have one or multiple logical drives per disk pool or array depending on the need. In addition, there is no limit on the size of the logical drive in an array. It can be as big as the array itself. However, in the current implementation of the disk pool, the maximum size of the logical drive that can be created in a disk pool is 64 TB. Contact IBM resellers or representatives to change this limit.

With controller firmware version 7.83 and later, on IBM DS Storage subsystem that can support disk pool, you can create a new type of logical drive that is referred to as 'thin volume' or 'thin logical drive'. Thin logical drive, which can be created with the actual physical disk allocation, is only a certain percentage of the size of the logical drive. The logical drive is presented to the host as a disk with full capacity. The actual physical disk allocation will grow in multiples of 4 GB allocations as needed. A usage level alert displays when the repository usage reaches a certain level. The user can then configure additional capacities as required. You can create a standard logical drive in a disk pool or an array. However, a thin logical drive can be created in a disk pool only. Thin provisioning is the ability of creating thin logical drives in the subsystem. This is provided free-of-charge with the required version of controller firmware on your IBM DS storage subsystem. Currently, only DS3500 and DCS3700 storage subsystems, and DCS3700 storage subsystems with Performance Module Controllers support thin provisioning.

When you create a standard logical drive, you allocate the available storage up front, configuring the capacity of a standard logical drive to meet application needs for data availability and I/O performance. To increase the storage available for host I/O data writes, you need to add additional drives to the disk pool or array. Thin logical drives let you create large virtual logical drives with small physical storage allocations that can grow over time to meet increased capacity demands. As storage demands increase, you can increase the amount of physical storage capacity as it is needed. Using thin logical drives helps to reduce the possibility of having excess, unused capacity in the disk pool.

Thin logical drives are logical drives that have a large virtual capacity available for host I/O data writes, but not all of the virtual capacity is associated with the allocated physical capacity. When you configure a thin logical drive, you specify two types of capacity: the virtual capacity and the physical/preferred capacity. The virtual capacity is the capacity that is reported to the host. The minimum value you can specify for the virtual capacity is 32 MB; the maximum is 64 TB. The physical capacity (also referred to as provisioned capacity and preferred capacity) is the amount of physical drive space that is currently allocated for writing data. An administrator can increase the physical capacity as capacity demands increase.

You can increase the physical capacity manually or automatically by specifying the maximum size of the physical capacity. The controller will automatically increase the physical capacity that it needs in 4 GB increments. If you select the manual method, the minimum value you can specify for the physical capacity is 4 GB; the maximum is 64 TB. You must specify the physical capacity in increments of 4 GB. IBM DS Storage Manager translates the specified physical capacity as a percentage of virtual capacity allocated for the thin logical drive. Typically, the allocated physical capacity is much smaller than the amount specified for the virtual capacity when the thin logical drive is created. The automatic increase of the physical capacity is the preferred method.

When the thin logical drive is created, an associated repository logical drive is also created for it. The initial size of the repository logical drive is the preferred capacity in the 'create logical drive' wizard. The default preferred capacity of the repository logical drive is 4 GB. An administrator can configure two attributes to help monitor capacity utilization of the repository logical drive and prevent a host write request from failing due to insufficient capacity. An administrator can set a repository logical drive utilization warning threshold percentage for the IBM DS Storage Manager to generate an alert when the specified percentage of capacity is utilized. To permit the subsystem to automatically expand the provisioned capacity by a specified amount when a repository utilization warning threshold is reached, an administrator specifies a maximum expansion value for the automatic expansion policy.

Some of the differences between a standard logical drive and a thin logical drive are:

- **RAID level** — The RAID level for a disk pool or an array determines how and where redundant data is distributed across the logical drives. The storage management software automatically assigns 10+P+Q RAID 6 for all the logical drives in a disk pool. For logical drives in an array, the storage management software lets you choose one of the supported RAID Supported logical drive types. You can create a standard logical drive in a disk pool or an array. You can create a thin logical drive in only a disk pool.
- **Capacity allocation** — With a standard logical drive you allocate the available storage capacity up front. With a thin logical drive you specify a virtual capacity and a preferred capacity, and increase the physical capacity to meet real capacity demands over time.
- **Capacity expansion** — You can increase the capacity of a standard logical drive in any increments. You must allocate capacity for a thin logical drive in increments of 4 GB because disk pool is create in multiples of 4 GB disk spaces.
- **Hot spare usage** — You can designate a hot spare drive for logical drives in an array for recovery from a drive failure condition. Hot spare drives are not used in disk pools. A percentage of each disk drive in a disk pool is reserved for reconstructed data in the event of a drive failure in a disk pool.

Creating a disk pool

To create a disk pool from unconfigured capacity nodes, complete the following steps in the Subsystem Management window:

Note: On a DCS3700 subsystem with Performance Module Controllers, 2 TB NL-SAS drives do not support T10 PI. Hence, you must manually create a disk pool if you are using these drives. This avoids T10 PI enabling. To manually create a disk pool, click **No** in the **Disk Pool Automatic Configuration** window.

1. To create a disk pool, complete one of the following:

- In the Storage & Copy Services tab, right-click **Unconfigured Capacity (HDD, SAS)** , and select **Create Disk pool**.
- In the Storage & Copy Services tab, select **Total Unconfigured Capacity**, and in the subsystem management window select **Storage > Disk pool > Create** .

The Create Disk Pool window opens.

2. Enter a new name for the Disk pool in **Disk pool name** and select the desired disk pool from the list.
3. Select **Create secured disk pool** if you want a secured disk pool.

Note: You can create a secured disk pool only if you have FDE drives and you have FDE premium feature enabled.

4. Click **View the notification settings**. You can view and change the critical warning and early warning notification thresholds.
5. Click **Create**. Wait till the disk pool is created.

Upon opening the **Subsystem Management** window or on deleting a logical drive, a disk pool, or an array, the **Disk Pool Automatic Configuration** window pops up, if -

- You have not created a disk pool and you have unassigned drives in the subsystem.
- You have created a disk pool and you have unassigned drives in the subsystem of the same type as in the disk pool.
- You have created multiple disk pools and you have unassigned drives in the subsystem of the same type as in those disk pools.

Adding capacity to an existing disk pool

You can add capacity to an existing disk pool. In the Storage & Copy Services tab, select the disk pool for which you want to add capacity and do one of the following:

- Select **Storage > Disk pool > Add Drives** in the subsystem management window.
- Right-click the disk pool to which you want to add drives, and select **Add Drives**.

Changing Disk Pool Settings

You can change critical warning and early warning notification thresholds, reconstruction priority, and disk pool preservation capacity of an existing disk pool. In the Storage & Copy Services tab, select the disk pool and do one of the following:

- Select **Storage > Disk pool > Change > Settings** in the subsystem management window.
- Right-click the disk pool of which you want to change settings and select **Change > Settings**.

The **Change Disk Pool Settings** window opens. You can increase or decrease the critical warning and early warning threshold limit, which is a percentage of the disk pool capacity. Slide the ruler to alter the Degraded reconstruction priority, Critical reconstruction priority, and background operation priority.

Click **OK** to save the settings.

Creating an array

To create an array from unconfigured capacity nodes, complete the following steps in the Subsystem Management window:

1. Use either of the following two methods to create a new array:
 - Select **Total Unconfigured Capacity**, and click **Array > Create**.
 - Select and right-click **Total Unconfigured Capacity**, and click **Create Array**. The Introduction (Create Array) window opens.
2. Click **Next**. The Array Name & Drive Selection (Create Array) window opens.

Note: If you are using 2 TB NL-SAS drives on a DCS3700 storage subsystem with Performance Module Controllers, select **Manual** for **Drive Selection**.

3. Take the applicable action for the following fields:
 - **Array name:** Enter a name for the new array. The name can be a maximum of 30 characters.
 - **Drive selection:** Select **Automatic** or **Manual (Advanced)**.

Automatic

Choose from a list of automatically generated drive and capacity options. This option is preselected by default.

Manual (Advanced)

Choose specific drives to obtain capacity for the new array.

- Click **Next**. The RAID Level and Capacity (Create Array) window opens.
4. Specify the RAID level (redundancy protection).
 5. Select the number of drives (overall capacity) for the new array.
 6. Click **Finish**. The Array Created window opens.
 7. If you want continue the process to create a logical drive, click **Yes**; if you want to wait to create a logical drive at another time, click **No**.

Redundant array of independent disks (RAID)

Redundant array of independent disks (RAID) is available on all operating systems and relies on a series of configurations, called *levels*, to determine how user and redundancy data is written and retrieved from the drives. The storage subsystem controller firmware supports six RAID level configurations:

- RAID-0
- RAID-1
- RAID-3
- RAID-5
- RAID-6
- RAID-10

Each level provides different performance and protection features. RAID-1, RAID-3, RAID-5, and RAID-6 write redundancy data to the drive media for fault tolerance. The redundancy data might be a copy of the data (mirrored) or an error-correcting code that is derived from the data. If a drive fails, the redundant data is stored on a different drive from the data that it protects. The redundant data is used to reconstruct the drive information on a hot-spare replacement drive. RAID-1 uses mirroring for redundancy. RAID-3, RAID-5, and RAID-6 use redundancy information, sometimes called *parity*, that is constructed from the data bytes and striped along with the data on each disk.

Table 17. RAID level descriptions

RAID level	Short description	Detailed description
RAID-0 Note: RAID-0 does not provide data redundancy.	Non-redundant, striping mode	RAID-0 offers simplicity, but does not provide data redundancy. A RAID-0 array spreads data across all drives in the array. This normally provides the best performance, but there is not any protection against single drive failure. If one drive in the array fails, all logical drives in the array fail. This RAID level must not be used for high data-availability needs. RAID-0 is better for noncritical data.
RAID-1 or RAID-10	Striping/ Mirroring mode	<ul style="list-style-type: none"> • A minimum of two drives are required for RAID-1: one for the user data and one for the mirrored data. The DS3000, DS4000, or DS5000 storage subsystem implementation of RAID-1 is a combination of RAID-1 and RAID-10, depending on the number of drives that are selected. If only two drives are selected, RAID-1 is implemented. If you select four or more drives (in multiples of two), RAID 10 is automatically configured across the array; two drives are dedicated to user data, and two drives are dedicated to the mirrored data. • RAID-1 provides high performance and the best data availability. On a RAID-1 logical drive, data is written to two duplicate disks simultaneously. On a RAID-10 logical drive, data is striped across mirrored pairs. • RAID-1 uses disk mirroring to make an exact copy of data from one drive to another drive. If one drive fails in a RAID-1 array, the mirrored drive takes over. • RAID-1 and RAID-10 are costly in terms of capacity. One-half of the drives are used for redundant data.
RAID-3	High-bandwidth mode	<ul style="list-style-type: none"> • RAID-3 requires one dedicated disk in the logical drive to hold redundancy information (parity). User data is striped across the remaining drives. • RAID-3 is a good choice for applications such as multimedia or medical imaging that write and read large amounts of sequential data. In these applications, the I/O size is large, and all drives operate in parallel to service a single request, delivering high I/O transfer rates.

Table 17. RAID level descriptions (continued)

RAID level	Short description	Detailed description
RAID-5	High I/O mode	<ul style="list-style-type: none"> RAID-5 stripes both user data and redundancy information (parity) across all of the drives in the logical drive. RAID-5 uses the equivalent of one drive capacity for redundancy information. RAID-5 is a good choice in multi-user environments such as database or file-system storage, where the I/O size is small and there is a high proportion of read activity. When the I/O size is small and the segment size is appropriately chosen, a single read request is retrieved from a single individual drive. The other drives are available to concurrently service other I/O read requests and deliver fast read I/O request rates.
RAID-6	Block-level striping with dual distributed parity	<p>RAID-6 is an evolution of RAID-5 and is designed for tolerating two simultaneous disk drive failures by storing two sets of distributed parities:</p> <ul style="list-style-type: none"> RAID Level 6 uses the equivalent of the capacity of two drives (in an array) for redundancy data. RAID Level 6 protects against the simultaneous failure of two drives by storing two sets of distributed parities. <p>Note: Not all DS storage subsystems support RAID-6. Check the announcement letter or the <i>Installation, User's, and Maintenance Guide</i> for your storage subsystem to determine whether RAID-6 and the minimum version of controller firmware required is supported for your storage subsystem.</p>

Note: One array uses a single RAID level, and all redundancy data for that array is stored within the array.

The capacity of the array is the aggregate capacity of the member drives, minus the capacity that is reserved for redundancy data. The amount of capacity that is needed for redundancy depends on the RAID level that is used.

To perform a redundancy check, click **Advanced > Recovery > Check array redundancy**. The redundancy check performs one of the following actions:

- Scans the blocks in a RAID-3, RAID-5, or RAID-6 logical drive and checks the redundancy information for each block
- Compares data blocks on RAID-1 mirrored drives

Important: When you select **Check array redundancy**, a warning message opens that informs you to use the option only when you are instructed to do so by the Recovery Guru. It also informs you that if you have to check redundancy for any reason other than recovery, you can enable redundancy checking through Media Scan.

Creating a standard logical drive

A *standard logical drive* is a logical structure that is the basic structure that you create to store data on the storage subsystem. The operating system recognizes a logical drive as a single drive. You can create a logical drive from an array or disk pool. Choose a RAID level to meet application needs for data availability and to maximize Fibre Channel I/O performance if you have created a logical drive from an array.

Note: For cluster configurations, if you add or delete logical drives, you must make them known to both nodes A and B.

To create a logical drive, complete the following steps in the Subsystem Management window:

1. On the Logical or Physical page of the Introduction (Create Logical Drive) window, click **Free Capacity** for an array or disk pool for which you want to create a new logical drive, right-click the array or disk pool, and click **Create Logical Drive**.
2. In the Specify Capacity/Name (Create Logical Drive) window, specify the following parameters for the logical drive that you are creating:

New logical drive capacity

The capacity can be the entire unconfigured capacity in an array or a portion of the array capacity.

Units Select GB, MB, or TB, depending upon the available capacity.

Name Type a name that is unique in the storage subsystem, up to a maximum of 30 characters.

3. In the “Advanced logical drive parameters” window, specify the applicable I/O characteristics (characteristics type, segment size, and cache read-ahead multiplier) and click **Next**. The Specify Logical Drive-to-LUN Mapping (Create Logical Drive) window opens.

Note: The I/O characteristics settings can be set automatically or they can be specified manually, according to one of the following logical drive usages: file system, database, or multimedia.

4. In the Specify Logical Drive-to-LUN Mapping (Create Logical Drive) window, specify the logical drive-to-LUN mapping.

The logical drive-to-LUN mapping preference can be one of the following two settings:

Default mapping

The Automatic setting specifies that a LUN is automatically assigned to the logical drive, using the next available LUN within the default host group. This setting grants logical drive access to host groups or host computers that have no specific logical drive-to-LUN mappings (those that were designated by the default host group node in the Topology view). If the Storage Partition feature is not enabled, you must specify the Automatic setting. In addition, you can also change the host type to match the host operating system.

Map later using the Mappings View

This setting specifies that you are not going to assign a LUN to the logical drive during creation. This setting enables you to define a specific logical drive-to-LUN mapping and create storage partitions, using the **Mappings Defined** option. Specify this setting when you enable storage partitioning.

5. Click **Finish** to create the logical drive. The Creation Successful (Create Logical Drive) window opens.
6. If you want to create another logical drive, click **Yes** in the Creation Successful (Create Logical Drive) window and proceed to step 9; otherwise, click **No**. When the Completed (Create Logical Drive) window opens, click **OK**, and continue with step 10.
7. In the Allocate Capacity (Create Logical Drive) window, choose to create the new logical drive from free capacity on the same array, free capacity on a different array, or from unconfigured capacity (create a new array). Repeat the process, beginning with step 1. The Completed (Create Logical Drive) window opens.
8. Click **OK**.
9. Register the logical drive with the operating system.

After you create logical drives with automatic logical drive-to-LUN mappings, follow the applicable instructions for your operating system in “Identifying devices” on page 159 to discover the new logical drive.

Creating a thin logical drive

To create a thin logical drive on a DS3500 or DCS3700 storage subsystem, or a DCS3700 storage subsystem with Performance Module Controllers having controller firmware 7.8x.xx.xx or later:

1. Do one of the following:
 - Right-click a Free Capacity node in the disk pool and select **Create Logical Drive**.
 - Select a Free Capacity node in the disk pool and select **Storage > Create > Logical Drive** from the main menu.

A **Create Logical Drive:Specify Parameters** window opens.

2. After giving it a name and specifying initial capacity, select the **Create thin logical drive** check box to create a thin logical drive.

Note: In a thin logical drive, dynamic read cache is not available.

3. Click **Next**. A **Create Logical Drive:Choose Physical capacity** window opens.
4. To create a thin logical drive with default settings, click **Finish**. To change any of the default settings, select **Customize capacity settings(advanced)** and click **Next**. A **Create Logical Drive:Customize settings** window opens.
5. You can choose your own preferred expansion capacity, maximum expansion capacity, and utilization warning threshold. Click **Finish**.

You have created a thin logical drive.

About Dynamic Capacity Expansion

Dynamic Capacity Expansion (DCE) is a modification operation in the storage management software that increases the capacity of an array or a disk pool. This modification operation enables you to add unassigned drives to an array or disk pool. Adding unassigned drives increases the free capacity in the array or the disk pool. You can use this free capacity to create additional logical drives or add reserve capacity for a disk pool or an array. This operation is considered to be dynamic because you have the ability to continually access data in the array throughout the entire operation. Keep these guidelines in mind when you add unassigned drives to an array or a disk pool:

- The number of unassigned drives that you can select for a DCE modification operation is limited by the controller firmware. You can add two unassigned drives at a time for an array. You can add up to twelve drives at a time for a disk pool. However, after you have completed a DCE operation, you can add more drives again until the desired capacity is reached.
- The existing logical drives in the array or the disk pool do not increase in size when you add unassigned drives to expand the free capacity. This operation redistributes existing logical drive capacity over the larger number of drives in the array or the disk pool.
- The unassigned drives that you are adding to the array or disk pool must be of the same media type and interface type. Mixing different drive types within a single array or disk pool is not permitted. Whenever possible, select drives that have a capacity equal to the capacities of the current drives in the array or the disk pool. Drives with capacities larger or equal to the capacity of the drives currently in a disk pool may be added as part of a Dynamic Capacity Expansion operation. However, if drives with larger capacity than those currently in a disk pool are added, the capacity above the smallest drive in the disk pool is not used and the amount unused is reported as unusable capacity.
- In a RAID Level 1 array, you must add two drives to make sure that data redundancy is configured.
- Only security capable drives can be added to a security enabled array or disk pool, or a security capable array or disk pool.
- In an array or a disk pool that is T10PI-capable and contains a T10PI-enabled logical drive, you can add only T10PI-capable drives.

Viewing Operations in Progress

The Operations in Progress window displays all of the long-running operations that are currently running in the storage subsystem. This is a view-only window and you can monitor progress. You can do other tasks in the Subsystem Management Window while the Operations in Progress window is open.

You can view progress for the following

- **Dynamic Capacity Expansion (DCE)** – Adding capacity to an array
- **Dynamic RAID Migration (DRM)** – Changing the RAID level of an array
- Checking the data redundancy of an array
- Defragmenting an array
- Initializing a logical drive
- **Dynamic Logical Drive Expansion (DVE)** – Adding capacity to a logical drive
- **Dynamic Segment Size (DSS)** – Changing the segment size of a logical drive
- **Reconstruction** – Reconstructing data from parity because of unreadable sectors or a failed drive
- **Copyback** – Copying data from a hot spare drive to a new replacement drive
- VolumeCopy
- Synchronizing a remote mirror

Refer to the online help for more information.

Configuring global hot-spare drives

You can assign available physical drives in the storage subsystem as a global *hot-spare drives* to keep data available. A global hot spare contains no data and acts as a standby in case a drive fails in a RAID 1, RAID 10, RAID 3, RAID 5, or RAID 6 arrays. A global hot spare is not used as a spare for a disk pool. The disk pool is created with reserved capacity distributed across all drives in the disk pool for reconstructed data in the event of a failed drive. If a drive in an array fails, the controllers automatically use a hot-spare drive to replace the failed physical drive while the storage subsystem is operating. The controller uses redundancy data to automatically reconstruct the data from the failed physical drive to the replacement (hot-spare) drive. This is called *reconstruction*. The hot-spare drive adds another level of redundancy to the storage subsystem. If a physical drive fails in the storage subsystem, the hot-spare drive is automatically substituted without requiring a physical swap.

Assigning hot-spare drives

There are two ways to assign hot-spare drives for arrays defined in the storage subsystem:

- **Automatically assign drives:** If you select this option, hot-spare drives are automatically created for the best hot-spare coverage, using the drives that are available. This option is always available.
- **Manually assign individual drives:** If you select this option, hot-spare drives are created out of those drives that were previously selected.

If you choose to manually assign the hot-spare drives, select a drive with a capacity equal to or larger than the total capacity of the drive that you want to cover with the hot spare. For example, if you have an 18 GB drive with configured capacity of 8 GB, you could use a 9 GB or larger drive as a hot spare. Generally, you must not assign a drive as a hot spare unless its capacity is equal to or greater than the capacity of the largest drive in the storage subsystem. For maximum data protection, you must use only the largest capacity drives for hot-spare drives in mixed capacity hard drive configurations. There is also an option to manually unassign individual drives.

If a drive fails in the array, the hot spare can be substituted automatically for the failed drive without requiring your intervention. If a hot spare is available when a drive fails, the controller uses redundancy data to reconstruct the data onto the hot spare.

Note: Drives with different interface protocols or technologies cannot be used as hot-spares for each other. For example, SATA drives and Fibre Channel drives cannot act as hot spares for each other.

1. To assign hot-spare drives, complete one of the following steps:
 - In the Storage & Copy Services tab, right-click the storage subsystem and select **Configuration > Hot Spare Coverage**.
 - In the Setup tab, select **Create Storage > Configure hot spares**.
2. Select Automatic or Manual. If you select Automatic, the storage subsystem automatically assigns a drive as hot-spare. You can later view/change hot spare coverage.
3. To select hot-spare drives manually, select the array and click **Assign**.
4. Select drives as hot-spare and click **OK**.

Restoring data from hot-spare drives

After the failed drive is physically replaced, you can use either of the following options to restore the data:

- When you have replaced the failed drive, the data from the hot spare is copied back to the replacement drive. This action is called *copyback*.
- You can assign the hot spare as a permanent member of the array. Performing the copyback function is not required for this option.

Note: If the controller firmware version is 10.84 or later, the hot-spare drive works in a copyback mode.

If you do not have a hot spare, you can still replace a failed drive while the array is operating. If the drive is part of a RAID Level 1, RAID Level 3, RAID Level 5, RAID Level 6, or RAID Level 10 array, the controller uses redundancy data to automatically reconstruct the data onto the replacement drive.

If you select **Manually unassign drives**, the hot-spare drives that you selected on the Physical tab are unassigned. This option is not available if you have not selected any drives on the Physical tab.

Defining a default host operating system

Before you use the logical drives in a host computer, you must specify the correct host type. The host type determines how the storage subsystem controllers work with each operating system on the hosts to which they are connected. If all of the host computers that are connected to the same storage subsystem are running the same operating system and you do not want to define partitioning, you can define a default host type.

To verify the current default host type, complete the following steps:

1. In the Subsystem Management window, click **View Storage Subsystem Profile**. A Storage Subsystem Profile window opens.
2. Click the **Host Mappings** tab. Right-click **Default Group** and select **Change Default Host Operating System**. The host-type name of the index that has the word Base next to it is the default host type.
3. Click **Close**.

Note: To enable Asymmetric Logical Unit Access (ALUA) multipath functionality the ALUA host type must be selected for the host partition or the default host group.

The host-type setting that you specify when you configure Storage Manager determines how the storage subsystem controllers work with the operating systems on the connected hosts. All Fibre Channel HBA ports that are defined with the same host type are handled in the same way by the storage subsystem controllers. This determination is based on the specifications that are defined by the host type. Some of the specifications that differ according to the host-type setting include the following options:

Auto Drive Transfer/ALUA

Enables or disables the Auto-Logical Drive Transfer feature (ADT). Starting with controller firmware version 7.83.xx.xx, the setting is used to enable or disable ALUA functionality instead.

Allow Reservation on Unowned LUNs

Determines the controller response to Reservation/Release commands that are received for LUNs that are not owned by the controller.

Reporting of Deferred Errors

Determines how the storage subsystem controller deferred errors are reported to the host.

Do Not Report Vendor Unique Unit Attention as Check Condition

Determines whether the controller reports a vendor-unique Unit Attention condition as a Check Condition status.

World Wide Name In Standard Inquiry

Enables or disables Extended Standard Inquiry.

Ignore UTM LUN Ownership

Determines how an inquiry for the Universal Access LUN (UTM LUN) is reported. The UTM LUN is used by Storage Manager to communicate to the storage subsystem in in-band management configurations.

Report LUN Preferred Path in Standard Inquiry Data

Reports the LUN preferred path in bits 4 and 5 of the Standard Inquiry Data byte 6.

Enable Host support for T10PI

Enable or Disable Host support for T10PI. If disabled, the controller strips the additional 8 bytes with T10PI information before sending data to the host.

In most storage subsystem configurations, the NVSRAM settings for each supported host type for a particular operating-system environment are sufficient for connecting a host to the storage subsystem. You do not have to change any of the host type settings for NVSRAM. If you think you have to change the NVSRAM settings, contact your IBM support representative before you proceed.

To define a default host type, complete the following steps:

1. Click **Host Mappings > Default Group > Change Default Host Operating System**. The Default Host-type window opens.
2. From the list, select the host type.
3. Click **OK**.

Note:

In the Veritas Storage Foundation Linux environment, the default host type must be set to 13.

Host type VMWARE has been added to NVSRAM as an additional host type. DS4200 and DS4700 will use index 21.

All other supported systems will use index 16.

Although not required, if using a Linux host type for a VMWARE host, it is recommended to move to the VMWARE host type since any upgrading of controller firmware and NVSRAM would continue to require running scripts, whereas using the VMWARE host type does not require running scripts.

- The controllers do not need to be rebooted after the change of host type
- The host will need to be rebooted

- Changing the host type should be done under low I/O conditions

Defining a host group

A *host group* is an entity in the Storage Partitioning topology that defines a logical collection of host computers that require shared access to one or more logical drives. You can grant individual hosts in a defined host group access to storage partitions, independently of the host group. You can make logical drive-to-LUN mappings to the host group or to an individual host in a host group.

You must create the host group at the storage subsystem level; do not create host groups at the default group level. However, you can use the default host group if you are running a storage subsystem configuration without partitioning enabled.

To define a host group, complete the following steps:

1. Click the **Host Mappings** tab on the Subsystem Management window.
2. In the Topology section of the Host Mappings page, highlight the name of the storage subsystem or Default Group, and right-click **Define > Host Group**.

Note: Ensure that the storage subsystem is highlighted in the left pane of the Subsystem Management window. Do not highlight Undefined Mappings.

3. Type a name for the new host group. Click **Ok**.

Defining heterogeneous hosts

The heterogeneous hosts feature enables hosts that are running different operating systems to access a single storage subsystem. Storage Manager supports up to 512 storage partitions on some subsystems, which enables a multiple host-type storage subsystem to share storage capacity, consolidate storage, and reduce storage management costs.

Host computers can run on different operating systems or variants of the same operating system. When you define a host type in the Define New Host Port window, the heterogeneous hosts feature enables the controllers in the storage subsystem to tailor their behavior (such as LUN reporting and error conditions) to the needs of the operating system or variant of the host that is sending the information.

Note:

1. During host-port definition, you must set each host type to the applicable operating system so that the firmware on each controller can respond correctly to the host.
2. You must enable storage partitioning, which is a premium feature. You must use the partition key that you saved at installation or go to the IBM webpage for feature codes to reactivate and obtain a new feature key. For more information about premium features, see “Storage Manager premium features” on page 50.
3. All the HBA ports that are connected to the storage subsystem from a server must be defined in a single host partition.

Note: To enable Asymmetric Logical Unit Access (ALUA) multipath functionality the ALUA host type must be selected for the host partition or the default host group.

Defining the host and host ports

To define the host and host ports by using the Define a host and host ports wizard, complete the following steps:

1. In the Topology section of the Mappings view of the Subsystem Management window, right-click the new host group and select **Define Host**. The Introduction (Define Host) window opens.
2. Click **Next**. The Specify Host Port Identifiers (Define Host) window opens.
3. Select the desired host port interface in **Choose a host interface type**.

Note: If you are configuring storage for IBM i, use the port on the first adapter. IBM i requires two adapters for a valid configuration.

4. If the HBA connectivity is properly set up, select **Add by selecting a known unassociated host port identifier** to add a host port. To enter WWPNN manually, select **Add by creating a new host port identifier**.
5. Enter the Host Port Identifier in **Alias** and click **Add**. It gets added to the Host Group list. Click **Next**.
6. Repeat steps 3 and 4 to add all hosts.
7. To remove any of the Host Port Identifiers you added, select the Host Port Identifiers in **Host port identifiers to be associated with the host** and click **Remove**.
8. Select the Host type and click **Next**.
9. Select **Yes** or **No** for share access to the same logical drives with other hosts. Click **Next**.
10. Review the information for accuracy. Click **Finish** to complete the task. If you want to define another host, select **Yes** and repeat the procedure again.

Mapping LUNs

This section describes how to map LUNs to a storage partition with the following procedures:

- “Mapping LUNs to a new Host or Host Group”
- “Adding LUNs to an existing Host or Host Group” on page 85

Mapping LUNs to a new Host or Host Group

To map LUNs to a newly created partition, complete the following steps:

1. Select the Mappings view of the Subsystem Management window.
2. In the Topology section, right-click the host on which you want to map LUNs, and select **Define Storage Partitioning**. The Define Storage Partitioning window opens.
3. In the Define Storage Partitioning window, select **Host** or **Host Group** and click **Next**.
4. Select the logical drive.
5. Either accept the default LUN ID or change it, and click **Add**.
6. Repeat step 5 for each LUN that you want to map to the partition.

Note: You can also use the Storage Partitioning wizard feature of the Storage Manager Task Assistant to map LUNs to a new storage partition.

Adding LUNs to an existing Host or Host Group

To map new LUNs to an existing partition, complete the following steps. Repeat these steps for each LUN that you want to add to the partition.

1. Click the **Host Mappings** tab on the Subsystem Management window.
2. In the Topology section, right-click the host or host group on which you want to map LUNs, and select **Define Additional Mappings**. The Define Additional Mapping window opens.
3. In the Define Additional Mapping window, select the following options, and then click **Add**:
 - Host group or host
 - Logical unit number (LUN)(0-255)
 - Logical drive

Configuring the IBM System Storage for IBM i

Use the information in the following sections, in combination with the “Configuring disk storage” on page 67 and “Defining a host group” on page 83 sections, to set up and assign IBM i LUNs on the DS5100 and DS5300 storage subsystems with the Storage Manager software.

Assigning a port identifier for IBM i

When you use the Storage Manager to enter a port identifier for IBM i, the port is on the first adapter. IBM i requires two adapters to make a valid configuration. The following illustration shows the setup window where you assign the port identifier.

DS4700_DtL - 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 iSeries.

[How do I match a host port identifier to a host?](#)

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:

21:01:00:e0:8b:a3:ec:cc

Refresh

☐ Add by creating a new host port identifier

New host port identifier (16 characters required):

Alias (30 characters maximum):

testHostPort1

Add ? Remove ?

Host port identifiers to be associated with the host:

Host Port Identifier	Alias / User Label
----------------------	--------------------

< Back Next > Cancel Help

Figure 10. Assigning a port identifier for IBM i

Defining IBM i as the host type

When you use the Storage Manager to define a host type, select **IBM i** from the **Host type (operating system)** list.

Important: In advanced setups, a LUN 0 might be assigned to a host group or host definition that does not allow IBM i as the host type. To fix this problem, remove the LUN 0 setting, change the operating system to IBM i, and then add the LUN that you previously removed.

The following illustration shows the setup window where you define the IBM i as the host type.

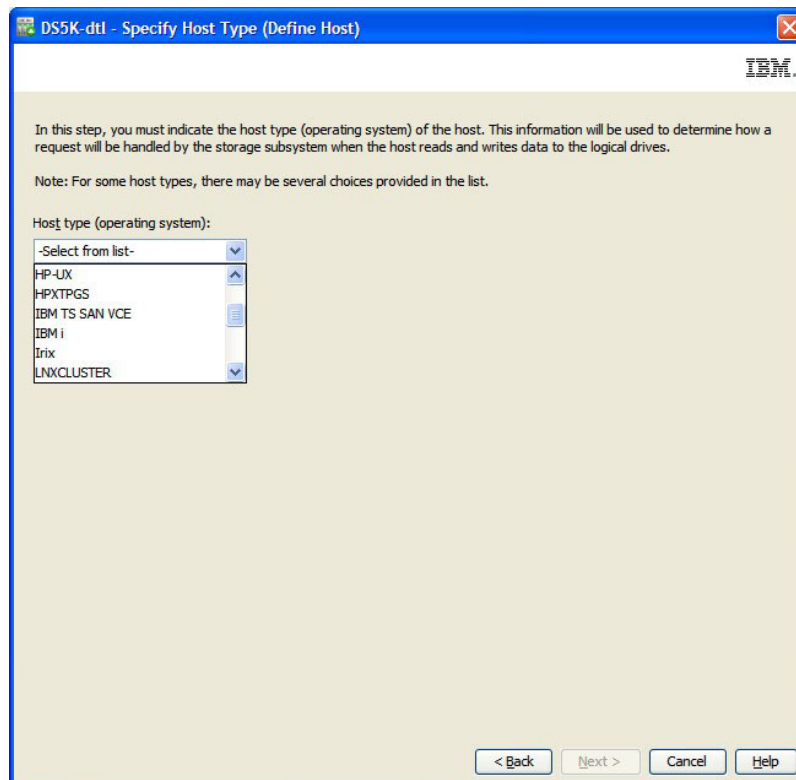


Figure 11. Selecting IBM i as the host type

Configuring and using optional premium features

This section describes optional premium features, including FlashCopy, Enhanced FlashCopy, VolumeCopy, Remote Mirror, and Full Disk Encryption.

Note: For more information about these optional premium features, see the *IBM System Storage DS Storage Manager Copy Services User's Guide* or contact your IBM reseller or IBM marketing representative.

About Enhanced FlashCopy

The DS Storage Manager with controller firmware version 7.8x.xx.xx or later, provides an Enhanced FlashCopy feature that lets you take a logical copy of the content of a standard logical drive or a thin logical drive at a particular point in time. These point-in-time images of logical drives are named Enhanced FlashCopy images. Taking Enhanced FlashCopy images are useful any time you need to be

able to roll back to a known good data set at a specific point in time. For example, you can create an Enhanced FlashCopy image as a backup that you can use during a recovery operation. Enhanced FlashCopy images of a base logical drive are managed as members of an Enhanced FlashCopy group. An Enhanced FlashCopy group can have up to 32 Enhanced FlashCopy images and a base logical drive allows creation of up to four Enhanced FlashCopy groups .

When you take an Enhanced FlashCopy image of a logical drive, the DS Storage Manager saves the Enhanced FlashCopy image in a repository associated with an Enhanced FlashCopy group. To provide a host access to Enhanced FlashCopy images stored within an Enhanced FlashCopy group repository, you must create an Enhanced FlashCopy logical drive of that Enhanced FlashCopy image.

Some important details to remember about the relationships between an Enhanced FlashCopy image, an Enhanced FlashCopy group, and an Enhanced FlashCopy logical drive are:

- Every Enhanced FlashCopy image is created in the context of exactly one Enhanced FlashCopy group.
- An Enhanced FlashCopy group is a sequence of Enhanced FlashCopy images of a single associated standard logical drive or thin logical drive. A logical drive used to create an Enhanced FlashCopy image is referred to as a base logical drive.
- An Enhanced FlashCopy group has exactly one repository that it uses to save the Enhanced FlashCopy images that are part of the Enhanced FlashCopy group
- All Enhanced FlashCopy images within an Enhanced FlashCopy group repository have a direct association with that Enhanced FlashCopy group.
- An Enhanced FlashCopy group has an association with a single logical drive.
- Every Enhanced FlashCopy logical drive has a direct association with an Enhanced FlashCopy image.
- Every Enhanced FlashCopy logical drive has a persistent relationship to the base logical drive of the Enhanced FlashCopy image for which the Enhanced FlashCopy logical drive was initially created.
- The repository associated with Enhanced FlashCopy logical drive has an association with an Enhanced FlashCopy group.

About FlashCopy

A FlashCopy logical drive is a logical point-in-time image of a logical drive, 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.

Note: The Storage subsystem can simultaneously utilize FlashCopy and Enhanced FlashCopy features. However, each base logical drive can use either of FlashCopy and Enhanced FlashCopy, but not both.

For additional information about the FlashCopy feature and how to manage FlashCopy logical drives, see the Storage Manager Subsystem Management window online help.

Important: The FlashCopy drive cannot be added or mapped to the same server that has the base logical drive of the FlashCopy logical drive in a Windows 2000, Windows Server 2003, or NetWare environment. You must map the FlashCopy logical drive to another server.

To create a FlashCopy logical drive, complete the following steps:

1. To make sure that you have the accurate point-in-time image of the base logical drive, stop applications and flush cache I/O to the base logical drive.
2. Open the Subsystem Management window. From the Logical page, right-click the base logical drive.
3. Select **Create FlashCopy Logical Drive**. The Create FlashCopy Logical Drive wizard starts.
4. Follow the on-screen instructions.
5. See the Subsystem Management window online help for instructions for adding the FlashCopy logical drive to the host.

Using VolumeCopy

The VolumeCopy feature is a firmware-based mechanism for replicating logical drive data within a storage subsystem. This feature is designed as a system-management tool for tasks such as relocating data to other drives for hardware upgrades or performance management, data backup, or restoring snapshot logical drive data. Users submit VolumeCopy requests by specifying two compatible drives. One drive is designated as the *source* and the other as the *target*. The VolumeCopy request is persistent so that any relevant result of the copy process can be communicated to the user. For more information about this feature, contact your IBM reseller or marketing representative.

Using Enhanced Remote Mirroring

Enhanced Remote Mirroring is a premium feature that is used for online, real-time replication of data between storage subsystems at different locations. Using the Remote Mirror option, you can designate a second storage subsystem to handle normal I/O operations if the first storage subsystem fails. For more information about this feature, see the *IBM System Storage DS Storage Manager Version 10 Copy Services User's Guide*, or contact the IBM reseller or marketing representative.

Using Enhanced Global Mirroring

Enhanced Global Mirroring is a premium feature that is used to replicate data between a local site and a remote site in an asynchronous scenario. This feature enables low cost fabric connectivity (iSCSI) and creates a temporary image on the primary logical drive. It periodically synchronizes with the secondary logical drive and thus minimizes system performance impact on non-high speed networks during peak hours. For more information about this feature, see the *IBM System Storage DS Storage Manager Version 10 Copy Services User's Guide*, or contact the IBM reseller or marketing representative.

Using Performance Read Cache

Performance Read Cache is a premium feature that is used to cache recently read data from the logical drives. This feature significantly increases data read

throughput on the cached data. Additional settings/administration are not required for its use. Performance Read Cache size depends on the installed controller cache size.

Note: Currently, this feature is supported on DS3500 and DCS3700 storage subsystems with Solid State Drive.

Table 18. Maximum supported Performance Read Cache size per installed controller cache

Installed controller cache size	Maximum supported Performance Read Cache size (per controller)
1 GB	1 TB
2 GB	2 TB
4 GB	4 TB
>4 GB	5 TB

Using Full Disk Encryption

Full Disk Encryption (FDE) is a premium feature that prevents unauthorized access to the data on a drive that is physically removed from a storage subsystem. Controllers in the storage subsystem have a security key. Secure drives provide access to data only through a controller that has the correct security key. FDE is a premium feature of the storage management software and must be enabled by you or your storage vendor.

Note: Not all DS storage subsystems support FDE. Check the announcement letter or the *Installation, User's, and Maintenance Guide* for your storage subsystem to determine whether FDE is supported for your storage subsystem.

The FDE premium feature requires security-capable drives. A security-capable drive encrypts data during write operations and decrypts data during read operations. Each security-capable drive has a unique drive-encryption key.

When you create a secure array from security-capable drives, the drives in that array become security-enabled. When a security-capable drive has been security-enabled, the drive requires the correct security key to read or write the data. A security-capable drive works the same as any other drive until it is security-enabled.

For more information about how to configure and use FDE, see Chapter 6, "Working with full disk encryption," on page 191.

Using other features

This section describes other features that are available in Storage Manager.

Using controller cache memory

Write caching enables the controller cache memory to store write operations from the host computer and, as a result, improves system performance. However, a controller can fail with user data in its cache that has not been transferred to the logical drive. Also, the cache memory can fail while it contains unwritten data. Write-cache mirroring protects the system from either of these possibilities. Write-cache mirroring enables cached data to be mirrored across two redundant controllers with the same cache size. The data that is written to the cache memory

of one controller is also written to the cache memory of the other controller. That is, if one controller fails, the other controller completes all outstanding write operations.

Note: You can enable the write-cache mirroring parameter for each logical drive, but when write-cache mirroring is enabled, half of the total cache size in each controller is reserved for mirroring the cache data from the other controller.

To prevent data loss or damage, the controller periodically writes cache data to the logical drive. When the cache holds a specified start percentage of unwritten data, the controller writes the cache data to the logical drive. When the cache is flushed down to a specified stop percentage, the flush is stopped. For example, the default start setting for a logical drive is 80% of the total cache size, and the stop setting is 20%. With these settings, the controller starts flushing the cache data when the cache reaches 80% full and stops flushing cache data when the cache is flushed down to 20% full.

For maximum data safety, you can choose low start and stop percentages, for example, a start setting of 25% and a stop setting of 0%. However, low start and stop settings increase the chance that data that is needed for a host computer read will not be in the cache. If sufficient data is not in the cache, the cache-hit percentage decreases, and subsequently, the I/O request rate decreases. It also increases the number of disk writes that are necessary to maintain the cache level, increasing system overhead and further decreasing performance.

If a power outage occurs, data in the cache that is not written to the logical drive might be lost, even if it is mirrored to the cache memory of both controllers. There are backup batteries in the controller enclosure that protect the cache against power outages.

Note: The controller backup battery CRU change interval is three years from the date that the backup battery CRU was installed for all models of the following DS4000 Storage Subsystems only: DS4100, DS4300, and DS4400. There is no replacement interval for the cache battery backup CRU in other DS4000 Storage Subsystems.

The Storage Manager software features a battery-age clock that you can set when you replace a battery. This clock keeps track of the age of the battery (in days) so that you know when it is time to replace the battery.

Note:

1. For the DS4100, and DS4300 or DS4300 Turbo disk systems, the battery CRU is inside each controller CRU.
2. For the DS4800, DS5100, and DS5300, the batteries CRU are in the interconnect-batteries CRU. Write caching is disabled when batteries are low or discharged. If you enable a parameter called write-caching without batteries on a logical drive, write caching continues even when the batteries in the controller enclosure are removed.

Attention: For maximum data integrity, do not enable the write-caching without batteries parameter, because data in the cache is lost during a power outage if the controller enclosure does not have working batteries. Instead, contact IBM service to get a battery replacement as soon as possible to minimize the time that the storage subsystem is operating with write-caching disabled.

Using Persistent Reservations

Attention: Use the Persistent Reservations option only with guidance from an IBM technical-support representative.

Use the Persistent Reservations option to view and clear logical drive reservations and associated registrations. Persistent reservations are configured and managed through the cluster server software and prevent other hosts from accessing particular logical drives.

Unlike other types of reservations, a persistent reservation is used to perform the following functions:

- Reserve access across multiple host ports and provide various levels of access control
- Query the storage subsystem about registered ports and reservations
- Provide for persistence of reservations in the event of a storage system power loss

You can use the Storage Manager software to manage persistent reservations in the Subsystem Management window. You can use the Persistent Reservation option to perform the following tasks:

- View registration and reservation information for all logical drives in the storage subsystem
- Save detailed information about logical drive reservations and registrations
- Clear all registrations and reservations for a single logical drive or for all logical drives in the storage subsystem

For detailed procedures, see the Subsystem Management window online help. You can also manage persistent reservations through the script engine and the command-line interface. For more information, see the Enterprise Management window online help.

Using Media Scan

A *media scan* is a background process that runs on all logical drives in the storage subsystem for which it is enabled, providing error detection on the drive media. The Media Scan feature checks the physical disks for defects by reading the raw data from the disk and, if there are errors, writing it back. The advantage of enabling Media Scan is that the process can find media errors before they disrupt normal logical-drive read and write functions. The media-scan process scans all logical-drive data to verify that it is accessible.

Note: The background media-scan operation does not scan hot-spare drives or unused optimal hard drives (those that are not part of a defined logical drive). To perform a media-scan operation on hot spares or unused optimal hard drives, you must convert them to logical drives at certain scheduled intervals and then revert them back to their hot-spare or unused states after you scan them.

There are two ways in which a media scan can run:

Logical drive redundancy checks not enabled

If background Media Scan is enabled with logical drive redundancy data checks not enabled, the storage subsystem scans all blocks in the logical drives, including the redundancy blocks, but it does not check for the accuracy of the redundancy data.

This is the default setting when you use Storage Manager to create logical drives.

Logical drive redundancy checks enabled

If background Media Scan is enabled with logical drive redundancy data checks enabled for RAID-3, RAID-5, or RAID-6 logical drives, a redundancy data check scans the data blocks, calculates the redundancy data, and compares it to the read redundancy information for each block. It then repairs any redundancy errors, if required. For a RAID-1 logical drive, a redundancy data check compares data blocks on mirrored drives and corrects any data inconsistencies.

Do not use this setting on older DS storage subsystems such as the DS4500, DS4400, DS4300, or DS4100; redundancy checking has a negative effect on storage subsystem performance.

For newer storage subsystems, such as the DS5100, DS5300, DS5020, or DS3950, this setting does not cause performance degradation.

When it is enabled, the media scan runs on all of the logical drives in the storage subsystem that meet the following conditions:

- The logical drive is in an optimal status.
- There are no modification operations in progress.
- The Media Scan parameter is enabled.

Note: 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.

Media Scan reads only data stripes, unless there is a problem. When a block in the stripe cannot be read, the read command is retried a certain number of times. If the read continues to fail, the controller calculates what that block must be and issues a write-with-verify command on the stripe. As the disk attempts to complete the write command, if the block cannot be written, the drive reallocates sectors until the data can be written. Then the drive reports a successful write and Media Scan checks it with another read. There must not be any additional problems with the stripe. If there are additional problems, the process repeats until there is a successful write, or until the drive is failed because of many consecutive write failures and a hot-spare drive takes over. Repairs are made only on successful writes, and the drives are responsible for the repairs. The controller issues only write-with-verify commands. Therefore, data stripes can be read repeatedly and report bad sectors, but the controller calculates the missing information with RAID.

In a dual-controller storage subsystem, there are two controllers that handle I/O (Controllers A and B). Each logical drive that you create has a preferred controller that normally handles I/O for it. If a controller fails, the I/O for logical drives that is *owned* by the failed controller fails over to the other controller. Media Scan I/O is not impacted by a controller failure, and scanning continues on all applicable logical drives when there is only one remaining active controller.

If a drive is failed during the media-scan process because of errors, normal reconstruction tasks are initiated in the controller operating system, and Media Scan attempts to rebuild the array using a hot-spare drive. While this reconstruction process occurs, no more media-scan processing is done on that array.

Note: Because additional I/O reads are generated for media scanning, there might be a performance impact, depending on the following factors:

- The amount of configured storage capacity in the storage subsystem. The greater the amount of configured storage capacity in the storage subsystem, the greater the performance impact.
- The configured scan duration for the media-scan operations. The longer the scan, the lower the performance impact.
- The status of the redundancy check option (enabled or disabled). If redundancy check is enabled, the performance impact is greater.

Errors reported by Media Scan

The media-scan process runs continuously in the background when it is enabled. Every time a media scan of all logical drives in a storage subsystem is completed, it restarts immediately. The media-scan process discovers any errors and reports them to the storage subsystem major event log (MEL). The following table lists the errors that are discovered during a media scan.

Table 19. Errors discovered during a media scan

Error	Description	Result
Unrecovered media error	The drive could not read the data on its first attempt or on any subsequent attempts.	For logical drives or arrays with redundancy protection (RAID-1, RAID-3 and RAID-5), data is reconstructed, rewritten to the drive, and verified. The error is reported to the event log. For logical drives or arrays without redundancy protection (RAID-0 and degraded RAID-1, RAID-3, RAID-5, and RAID-6 logical drives), the error is not corrected but is reported to the event log.
Recovered media error	The drive could not read the requested data on its first attempt but succeeded on a subsequent attempt. Note: Media scan makes three attempts to read the bad blocks.	The data is rewritten to the drive and verified. The error is reported to the event log.
Redundancy mismatches	Redundancy errors are found. Note: This error can occur only when the optional redundancy check box is selected, when the Media Scan feature is enabled, and the logical drive or array is not RAID-0.	The first 10 redundancy mismatches that are found on a logical drive are reported to the event log.
Unfixable error	The data could not be read, and parity or redundancy information could not be used to regenerate it. For example, redundancy information cannot be used to reconstruct data on a degraded logical drive.	The error is reported to the event log.

Media Scan settings

To maximize the protection and minimize the I/O performance impact, the storage subsystem comes from the manufacturer with the following default Media Scan settings:

- The Media Scan option is enabled for all logical drives in the storage subsystem. Therefore, every time a logical drive is created, it is created with the Media Scan option enabled. If you want to disable media scanning, you must disable it manually for each logical drive.
- The media-scan duration is set to 30 days. This is the time in which the storage subsystem controllers must complete the media scan of a logical drive. The controller uses the media-scan duration, with the information about which logical drives must be scanned, to determine a constant rate at which to perform the media-scan activities. The media-scan duration is maintained regardless of host I/O activity.

Thirty days is the maximum duration setting. You must manually change this value if you want to scan the media more frequently. This setting is applied to all logical drives in the storage subsystem. For example, you cannot set the media-scan duration for one logical drive at two days and the duration for other logical drives at 30 days.

- By default, the redundancy check option is not enabled on controller firmware versions earlier than 7.60.39.00. For controller firmware versions earlier than 7.60.39.00, you must manually set this option for each of the logical drives on which you want to have redundancy data checked.

For controller firmware version 7.60.39.00 and later, the redundancy check option is enabled as a default setting for any newly created logical drives. If you want an existing logical drive that was created before version 7.60.39.00 or later was installed to have the redundancy check option enabled, you must enable the option manually.

Without redundancy check enabled, the controller reads the data stripe to confirm that all the data can be read. If it reads all the data, it discards the data and moves to the next stripe. When it cannot read a block of data, it reconstructs the data from the remaining blocks and the parity block and issues a write with verify to the block that could not be read. If the block has no data errors, Media Scan takes the updated information and verifies that the block was fixed. If the block cannot be rewritten, the drive allocates another block to take the data. When the data is successfully written, the controller verifies that the block is fixed and moves to the next stripe.

Note: With redundancy check, Media Scan goes through the same process as without redundancy check, but, in addition, the parity block is recalculated and verified. If the parity has data errors, the parity is rewritten. The recalculation and comparison of the parity data requires additional I/O, which can affect performance.

Important: Changes to the Media Scan settings do not go into effect until the current media-scan cycle is completed.

To change the Media Scan settings for an entire storage subsystem, complete the following steps:

1. Select the storage subsystem entry on the Logical or Physical tab of the Subsystem Management window.
2. Click **Storage Subsystem > Change > Media Scan Settings**.

To change the Media Scan settings for a logical drive, complete the following steps:

1. Select the logical drive entry on the **Logical** or **Physical** tab of the Subsystem Management window.
2. Click **Storage Subsystem > Change > Media Scan Settings**.

Media Scan duration

When Media Scan is enabled, a duration window is specified (in days) which indicates how long the storage subsystem will give the media-scan process to check all applicable logical drives. The duration window can be shortened or increased to meet the customer requirements. The shorter the duration, the more often a drive is scanned and consequently, the more robust the situation will be. However, the more often a drive is scanned, the higher the performance impact.

Whenever the storage subsystem has some idle time, it starts or continues media scanning operations. If application generated disk I/O work is received, it gets priority. Therefore, the media-scan process can slow down, speed up, or in some cases be suspended as the work demands change. If a storage subsystem receives a great deal of application-generated disk I/O, it is possible for the Media Scan to fall behind in its scanning. As the storage subsystem gets closer to the end of the duration window during which it must finish the media scan, the background application starts to increase in priority (i.e. more time is dedicated to the media-scan process). This increase in priority only increases to a certain point because the storage subsystem priority is process application-generated disk I/O. In this case, it is possible that the media-scan duration will be longer than the media scan duration settings.

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

Tuning storage subsystems

The information in this section describes Performance Monitor data and the tuning options that are available in the Storage Manager to optimize storage subsystem and application performance. Use the Subsystem Management window Performance Monitor to monitor storage subsystem performance in real time and to save performance data to a file for later analysis. You can specify the logical drives and controllers to monitor and the polling interval. You can also receive storage subsystem totals, which is data that combines the statistics for both controllers in an active-active controller pair.

Table 20. Performance Monitor tuning options in the Subsystem Management window

Data field	Description
Total I/Os	Total I/Os that have been performed by this device since the beginning of the polling session.
Read percentage	The percentage of total I/Os that are read operations for this device. Write percentage is calculated as 100 minus this value.
Cache-hit percentage	The percentage of read operations that are processed with data from the cache, rather than requiring a read from the logical drive.
Current® KB per second	During the polling interval, the <i>transfer rate</i> is the amount of data, in KB, that is moved through the Fibre Channel I/O path in 1 second (also called throughput).
Maximum KB per second	The maximum transfer rate that is achieved during the Performance Monitor polling session.

Table 20. Performance Monitor tuning options in the Subsystem Management window (continued)

Data field	Description
Current I/O per second	The average number of I/O requests that are serviced per second during the current polling interval (also called an I/O request rate).
Maximum I/O per second	The maximum number of I/O requests that are serviced during a 1-second interval over the entire polling session.

Maximizing throughput with load balancing

Load balancing is the redistribution of read or write requests to maximize throughput between the server and the storage subsystem. Load balancing is very important in high-workload settings or other settings where consistent service levels are critical. The multipath driver balances I/O workload transparently, without administrator intervention. Without multipath software, a server that sends I/O requests down several paths might operate with heavy workloads on some paths while other paths are not used efficiently.

The multipath driver determines which paths to a device are in an active state and can be used for load balancing. The load balancing policy uses one of three algorithms: *round robin*, *least queue depth*, or *least path weight*. Multiple options for setting the load balance policies enable you to optimize I/O performance when mixed-host interfaces are configured. The load balancing policies that you can choose depend on your operating system. Load balancing is performed on multiple paths to the same controller, but not across both controllers.

Table 21. Load balancing policies supported by operating systems

Operating system	Multi-path driver	Load balancing policy
AIX	MPIO	Round robin, selectable path priority
Red Hat Enterprise Linux 4 Update 7	RDAC	Round robin, least queue depth
Solaris	MPxIO	Round robin
SUSE Linux Enterprise 9 Service Pack 4	RDAC	Round robin, least queue depth
Windows	MPIO	Round robin, least queue depth, least path weight

Round robin with subset

The round robin with subset I/O load balance policy routes I/O requests, in rotation, to each available data path to the controller that owns the logical drives. This policy treats all paths to the controller that owns the logical drive equally for I/O activity. Paths to the secondary controller are ignored until ownership changes. The basic assumption for the round robin policy is that the data paths are equal. With mixed host support, the data paths might have different bandwidths or different data transfer speeds.

Least queue depth with subset

The least queue depth with subset policy is also known as the least I/Os or least requests policy. This policy routes the next I/O request to a data path that has the

fewest queued outstanding I/O requests. For this policy, an I/O request is simply a command in the queue. The type of command or the number of blocks that are associated with the command are not considered. The least queue depth with subset policy treats large block requests and small block requests equally. The selected data path is one of the paths in the path group of the controller that owns the logical drive.

Least path weight with subset

The least path weight with subset policy assigns a weight factor to each data path to a logical drive. An I/O request is routed to the path with the lowest weight value to the controller that owns the logical drive. If more than one data path to the logical drive has the same weight value, the round-robin with subset path selection policy is used to route I/O requests between the paths with the same weight value.

Balancing the Fibre Channel I/O load

The **Total I/O** data field in the Subsystem Management window is used for monitoring the Fibre Channel I/O activity to a specific controller and a specific logical drive. This field helps you to identify possible I/O hot spots.

You can identify Fibre Channel I/O patterns to the individual logical drives and compare those with the expectations according to the application. If a controller has more I/O activity than expected, click **Array > Change Ownership** to move an array to the other controller in the storage subsystem.

It is difficult to balance Fibre Channel I/O loads across controllers and logical drives because I/O loads are constantly changing. The logical drives and the data that is accessed during the polling session depend on which applications and users are active during that time period. It is important to monitor performance during different time periods and gather data at regular intervals to identify performance trends. The Performance Monitor enables you to save data to a comma-delimited text file that you can import to a spreadsheet for further analysis.

If you notice that the workload across the storage subsystem (total Fibre Channel I/O statistic) continues to increase over time while application performance decreases, you might have to add storage subsystems to the enterprise.

Optimizing the I/O transfer rate

The transfer rates of the controller are determined by the application I/O size and the I/O request rate. A small application I/O request size results in a lower transfer rate but provides a faster I/O request rate and a shorter response time. With larger application I/O request sizes, higher throughput rates are possible. Understanding the application I/O patterns will help you optimize the maximum I/O transfer rates that are possible for a given storage subsystem.

One of the ways to improve the I/O transfer rate is to improve the I/O request rate. Use the host-computer operating system utilities to gather data about I/O size to understand the maximum possible transfer rates. Then, use the tuning options that are available in the Storage Manager to optimize the I/O request rate to reach the maximum possible transfer rate.

Optimizing the I/O request rate

The I/O request rate can be affected by the following factors:

- The I/O access pattern (random or sequential) and I/O size
- The status of write-caching (enabled or disabled)
- The cache-hit percentage
- The RAID level
- The logical-drive modification priority
- The segment size
- The number of logical drives in the arrays or storage subsystem
- The fragmentation of files

Note: Fragmentation affects logical drives with sequential I/O access patterns, not random I/O access patterns.

Determining the I/O access pattern and I/O size

To determine whether the I/O access has sequential characteristics, enable a conservative cache read-ahead multiplier (for example, 4) by clicking **Logical Drive > Properties**. Then, examine the logical drive cache-hit percentage to see whether it has improved. An improvement indicates that the I/O has a sequential pattern. Use the host-computer operating-system utilities to determine the typical I/O size for a logical drive.

Enabling write-caching

Higher I/O write rates occur when write-caching is enabled, especially for sequential I/O access patterns. Regardless of the I/O access pattern, be sure to enable write-caching to maximize the I/O rate and shorten the application response time.

Optimizing the cache-hit percentage

A higher cache-hit percentage is preferred for optimal application performance and is positively correlated with the Fibre Channel I/O request rate.

If the cache-hit percentage of all logical drives is low or trending downward and less than the maximum amount of controller cache memory is installed, you might have to install more memory.

If an individual logical drive has a low cache-hit percentage, you can enable cache read-ahead for that logical drive. Cache read-ahead can increase the cache-hit percentage for a sequential I/O workload. If cache read-ahead is enabled, the cache fetches more data, usually from adjacent data blocks on the drive. In addition to the requested data, this feature increases the chance that a future request for data is fulfilled from the cache, rather than requiring a logical drive access.

The cache read-ahead multiplier values specify the multiplier to use for determining how many additional data blocks are read into the cache. Choosing a higher cache read-ahead multiplier can increase the cache-hit percentage.

If you determine that the Fibre Channel I/O access pattern has sequential characteristics, set an aggressive cache read-ahead multiplier (for example, 8). Then examine the logical-drive cache-hit percentage to see whether it has improved. Continue to customize logical-drive cache read-ahead to arrive at the optimal multiplier (for a random I/O pattern, the optimal multiplier is 0).

Choosing appropriate RAID levels

Use the read percentage for a logical drive to determine the application behavior. Applications with a high read percentage perform well with RAID-5 logical drives because of the outstanding read performance of the RAID-5 configuration.

Note: This is applicable for traditional array only. Disk pools are always created with RAID 6 and 8D+P+Q.

Applications with a low read percentage (write-intensive) do not perform as well on RAID-5 logical drives because of the way that a controller writes data and redundancy data to the drives in a RAID-5 logical drive. If there is a low percentage of read activity relative to write activity, you can change the RAID level of a logical drive from RAID-5 to RAID-1 for faster performance.

Choosing an optimal logical-drive modification priority setting

The modification priority defines how much processing time is allocated for logical-drive modification operations versus system performance. The higher the priority, the faster the logical-drive modification operations are completed, but the more slowly the system I/O access pattern is serviced.

Logical-drive modification operations include reconstruction, copyback, initialization, media scan, defragmentation, change of RAID level, and change of segment size. The modification priority is set for each logical drive, using a slider bar from the Logical Drive - Properties window. There are five relative settings on the reconstruction rate slider bar, ranging from Low to Highest. The actual speed of each setting is determined by the controller. Choose the Low setting to maximize the Fibre Channel I/O request rate. If the controller is idle (not servicing any I/O request rates) it ignores the individual logical-drive rate settings and processes logical-drive modification operations as fast as possible.

Choosing an optimal segment size for an array

A segment is the amount of data, in KB, that the controller writes on a single logical drive before it writes data on the next drive. A data block is 512 bytes of data and is the smallest unit of storage. The size of a segment determines how many data blocks it contains. For example, an 8 KB segment holds 16 data blocks, and a 64 KB segment holds 128 data blocks.

Note: For a Disk pool, the segment size is always 128 KB. The segment size in the DS Storage Manager is expressed in KB.

When you create a logical drive, the default segment size is a good choice for the expected logical-drive usage. To change the default segment size, click **Logical Drive > Change Segment Size**.

If the I/O size is larger than the segment size, increase the segment size to minimize the number of drives that are needed to satisfy an I/O request. This technique helps even more if you have random I/O access patterns. If you use a single logical drive for a single request, it leaves other logical drives available to simultaneously service other requests.

When you use the logical drive in a single-user, large I/O environment such as a multimedia application, storage performance is optimized when a single I/O request is serviced with a single array data stripe (which is the segment size multiplied by the number of logical drives in the array that are used for I/O requests). In this case, multiple logical drives are used for the same request, but each logical drive is accessed only once.

Defragmenting files to minimize disk access

Each time that you access a drive to read or write a file, it results in the movement of the read/write heads. Verify that the files on the logical drive are defragmented. When the files are defragmented, the data blocks that make up the files are next to each other, preventing extra read or write head movement when files are retrieved. Fragmented files decrease the performance of a logical drive with sequential I/O access patterns.

Using the Storage Manager command-line interface and Script Editor

This section describes the Storage Manager command-line interface and the Script Editor.

Storage Manager command-line interface

Attention: The command-line interface (CLI) does not have any mechanisms to prevent you from inadvertently making unwanted changes to the storage subsystem. The script commands are capable of damaging a configuration and causing loss of data access if not used correctly. To avoid damaging effects to your storage configuration or data, use the Storage Manager client graphical user interface (GUI) to manage your storage subsystem configurations.

The command-line interface is a software tool that enables you to configure and monitor storage subsystems using script commands. Using the CLI, you can run commands from an operating-system prompt, such as the Windows command prompt, a Linux operating-system console, or a Solaris operating-system console. You must install the IBM DS Storage Manager client to run the script commands either through the script window, which is invoked from the IBM DS Storage Manager client Enterprise window, or through the command-line interface using the SMcli program. The script command engine is automatically installed as part of the IBM DS Storage Manager client installation.

Each command performs a specific action that manages a storage subsystem or returns information about the status of a storage subsystem. You can enter individual commands, or you can run script files when you need to perform operations more than once. For example, you can run script files when you want to install the same configuration on several storage subsystems. With the CLI, you can 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. You can use the CLI both in installation sites and in development environments.

For more information about the Storage Manager CLI, see the *IBM System Storage DS3000, DS4000, and DS5000 Command Line Interface and Script Commands Programming Guide*.

Using the Script Editor

Instead of using the graphical user interface to perform storage subsystem management functions, a Script Editor window is provided for running scripted management commands. If the controller firmware version is 5.4x.xx.xx or earlier, some of the management functions in the graphical user interface are not available through script commands. Storage Manager 10.xx, in conjunction with controller firmware version 07.xx.xx.xx and later, provides full support of all management functions through SMcli commands.

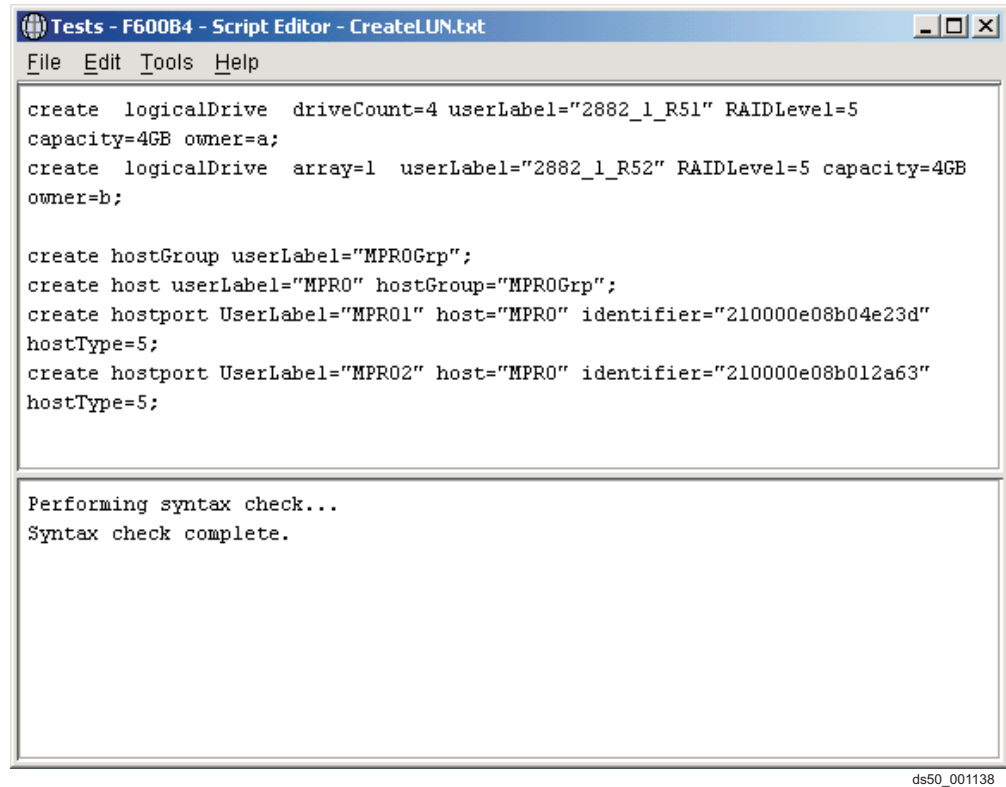


Figure 12. The Script Editor window

Important: Use caution when you run the commands; the Script Editor does not prompt you for confirmation of operations that are destructive, such as the **Delete arrays** and **Reset Storage Subsystem configuration** commands.

Not all script commands are implemented in all versions of the controller firmware. The earlier the firmware version, the smaller the set of available script commands. For more information about script commands and firmware versions, see the Storage Manager Enterprise Management window.

For a list of available commands and their syntax, see the online Command Reference help.

To open the Script Editor, complete the following steps:

1. Select a storage subsystem in either the tree view or the table view of the **Device** tab.
2. Click **Tools > Execute Script**.
3. The Script Editor opens. The Script view and the Output view are presented in the window. A splitter bar divides the window between the Script view and the Output view. Drag the splitter bar to resize the views.

In the Script view, you can input and edit script commands. The Output view displays the results of the operations. The Script view supports the following editing key strokes:

Ctrl+A

Selects everything in the window

Ctrl+C

Copies the marked text in the window into a Windows clipboard buffer

Ctrl+V

Pastes the text from the Windows clipboard buffer into the window

Ctrl+X Deletes (cuts) the marked text in the window

Ctrl+Home

Moves the cursor to the top of the script window

Ctrl+End

Moves the cursor to the bottom of the script window

The following list shows general guidelines for using the Script Editor:

- All statements must end with a semicolon (;).
- Each command and its associated primary and secondary parameters must be separated by a space.
- The Script Editor is not case-sensitive.
- Each new statement must begin on a separate line.
- Comments can be added to your scripts to make it easier for you and other users to understand the purpose of the command statements.

The Script Editor supports the two following comment formats:

- Text contained after two forward slashes (//) until an end-of-line character is reached

For example:

```
//The following command assigns hot spare drives.  
set drives [1,2 1,3] hotspare=true;
```

The comment `//The following command assigns hot spare drives.` is included for clarification and is not processed by the Script Editor.

Important: You must end a comment that begins with `//` with an end-of-line character, which you insert by pressing the Enter key. If the script engine does not find an end-of-line character in the script after processing a comment, an error message displays and the script fails.

- Text contained between the `/*` and `*/` characters

For example:

```
/* The following command assigns hot spare drives.*/  
set drives [1,2 1,3] hotspare=true;
```

The comment `/*The following command assigns hot spare drives.*/` is included for clarification and is not processed by the Script Editor.

Important: The comment must start with `/*` and end with `*/`. If the script engine does not find both a beginning and ending comment notation, an error message displays and the script fails.

Chapter 5. Configuring hosts

After you configure the storage subsystem or subsystems, use the information in this chapter to enable all hosts to connect to the storage subsystems. This chapter consists of the following sections:

- “Booting a host operating system using SAN boot”
- “Using multipath drivers to automatically manage logical drive fail-over and fail-back” on page 115
- “Identifying devices” on page 159
- “Configuring devices” on page 162

Booting a host operating system using SAN boot

SAN boot is the ability to boot the host operating system from a Storage Area Network (SAN) device. In this case, the device is a LUN from a DS3000, DS4000, or DS5000 storage subsystem. SAN boot is also referred to as *remote boot*, where the boot LUN is in the storage subsystem instead of inside the server enclosure and the server is connected to the storage subsystem in a SAN. The connections might be direct connections or through the SAN fabric - SAS, FC, or iSCSI.

Using SAN boot includes the following advantages:

Server consolidation

Each server can boot from an image of the operating system on the SAN.

Simplified recovery from server failures

Operating-system reinstallation is not required.

Rapid disaster recovery

The storage subsystem can be replicated at a remote recovery site.

The following conditions are required for SAN boot:

- SAN configuration, zoning of boot devices, multipath configurations (if applicable)
- Single active path to boot LUN. During the installation process, only one path to the boot LUN must be enabled before you install and enable a multipath driver.
- HBA BIOS; selectable boot, or boot BIOS, must be enabled.

To configure a storage subsystem for SAN boot, use the following guidelines:

1. Configure the SAN fabric:
 - a. Create SAN zoning and arrange the Fibre Channel devices into logical groups over the physical configuration of the fabric. Each device in a SAN might be placed into multiple zones.
 - b. Remove all paths from the server HBA ports except for one, to the boot LUN. To do this, disable the port on the switch for the other physical paths.
2. Configure the storage subsystem:
 - a. Create the LUN.
 - b. Map the LUN to the host as LUN 0.

Note: You must know the HBA WWNN, which you can get from the HBA label.

3. Configure the HBAs for boot from SAN:
 - a. Verify that boot BIOS is enabled on the HBA device configured for the host.
 - b. When the host is starting, enter the boot BIOS for your HBA device.
 - c. Select the HBA that you want to use for SAN booting and configure the BIOS so that the boot LUN is designated as the preferred boot device. After the storage subsystem has discovered the HBA WWPNNs, you must configure them as the HBAs to the boot LUN, using the host-mapping procedures.

Note:

- 1) The HBA must be logged in to the storage subsystem. Even though no LUN will be available yet, you can use the BIOS to discover the storage subsystem.
 - 2) For more information, see the documentation that came with your HBA.
 - d. Save the changes, exit BIOS, and restart the server. The BIOS can now be used to discover the newly configured LUN.
4. Start the installation by booting from the installation media:
 - a. During the installation, your operating-system media asks which drive (or LUN) you want to perform the installation. Select the drive that corresponds to your storage subsystem device.

Note: If you are prompted during the installation for third-party device drivers, select the HBA driver that you have available on another form of media.

- b. Choose the default option for disk partitioning.

Note: Make sure that the LUN you choose is large enough for the operating system. For Linux, and most other operating systems, 20 GB is enough for the boot device. For swap partitions, make sure that the size is at least the size of your server physical memory.

5. Complete the installation and finish the SAN boot procedure:
 - a. Restart the server again, and open the boot options menu. The boot device that you set up is ready to be used.
 - b. Select the option to boot from a hard disk drive/SAN, and select the HBA that is associated with the SAN disk device on which the installation was completed. The installation boot device is now listed in the bootable devices that are discovered on the selected HBA.
 - c. Select the applicable device, and boot.
 - d. Set the installed boot device as the default boot device for the system.

Note: This step is not required. However, the installed boot device must be the default boot device to enable unattended reboots after this procedure is complete.

- e. **Linux only – To complete the installation on Linux, complete the following steps:**

- 1) Verify that the persistent binding for `/var/mpp/devicemapping` is up-to-date. The `/var/mpp/devicemapping` file tells RDAC which storage subsystem to configure first. If additional storage subsystems will be added to the server, the storage subsystem with the boot/root volume must always be first in the device mapping file. To update this file, execute the following command:

```
# mppUpdate
```

- 2) After you run the `# mppUpdate` command, cat the `/var/mpp/devicemapping` file with the following command:

```
# cat /var/mpp/devicemapping 0:<DS4x00 SAN Boot Device>
```

The storage subsystem for the boot/root volume must be at entry 0. If the boot/root volume is not at entry 0, edit the file to reorder the storage subsystem entries so that the array for the boot/root volume is at entry 0.

- 3) Execute the `# mppUpdate` command. The installation is now complete.

Additional paths between the storage subsystem and server can now be added. If the server is going to be used to manage the storage subsystem, the Storage Manager can now be installed on the server.

For additional information about using multipath drivers, see “Using multipath drivers to automatically manage logical drive fail-over and fail-back” on page 115.

Overview of multipath drivers

One of the primary functions of the **failover** feature is to provide path management. If there is more than one path from the server to the controller, some multipath drivers are also able to spread the I/Os (input/output) between the paths. Please check the documentation with the multipath failover driver for this support.

Note: The connections between the hosts and the storage subsystems in the following figures are meant to illustrate the concept of multipath drivers. These are not recommendations.

Figure 13 on page 106 and Figure 14 on page 107 show how the I/Os flow in the optimal single and two paths from server to controller environment.

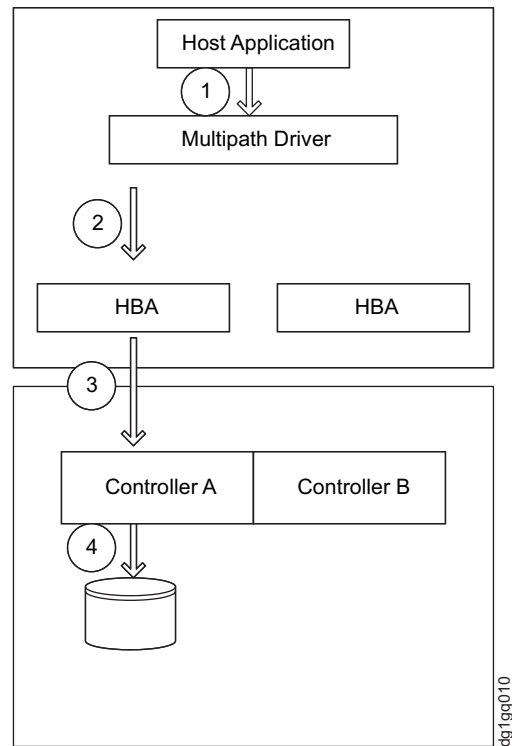


Figure 13. I/O flow in an optimal single path

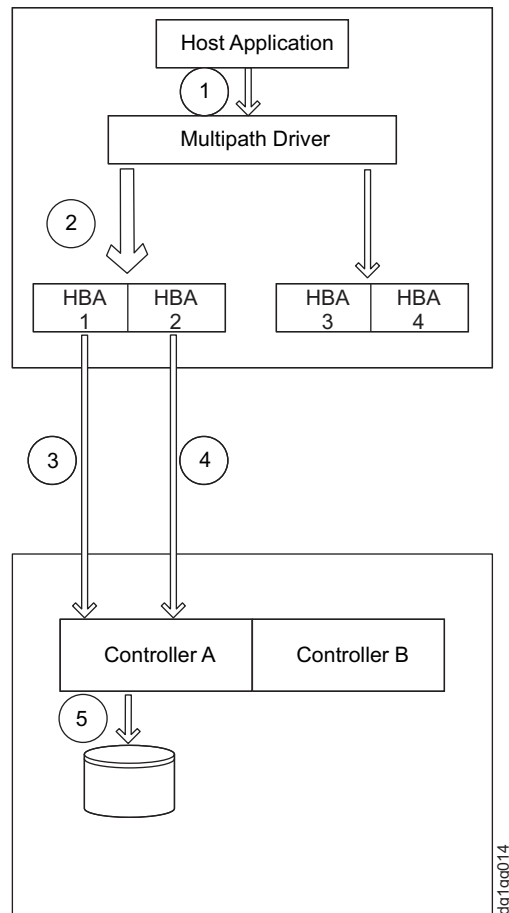


Figure 14. I/O flow in optimal two path

Figure 14 also illustrates that I/Os to the logical drive can be round robin through all of the available paths if the multipath driver supports.

Failover

The multipath drivers monitor the data path to the storage subsystem if they are not working correctly or in case of multiple link errors. If multipath drivers detect either of these conditions, it checks the path table for the redundant paths and controller. The failover driver performs a path failover if alternate paths to the same controller are available. Figure 15 on page 108 shows that the multipath driver uses only one of the two paths to the controller because the other path fails. If all of the paths to a controller fail, the multipath driver performs a controller failover as shown in Figure 16 on page 109 and Figure 17 on page 110. Here when controller A fails, the multipath driver moves the ownership of the logical drive from controller A to controller B. The controller B then receives and process all I/Os to the logical drive.

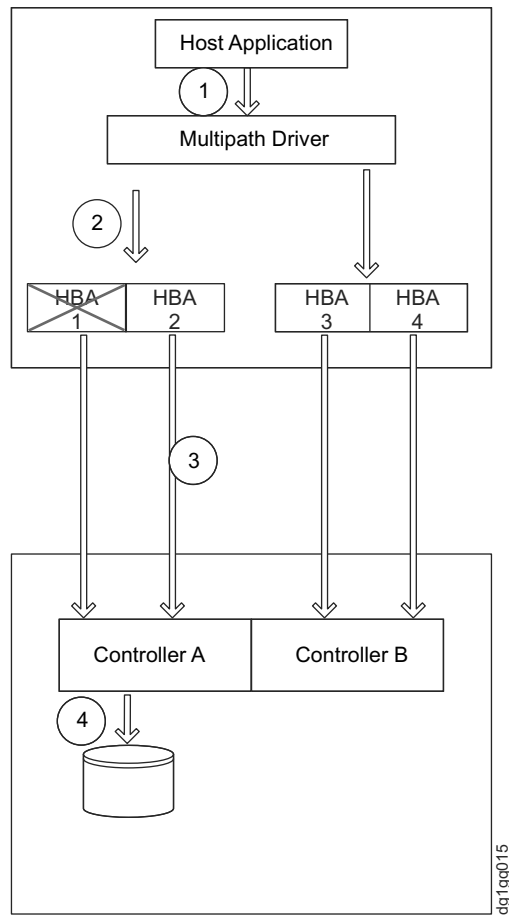


Figure 15. Use of one path when the other fails.

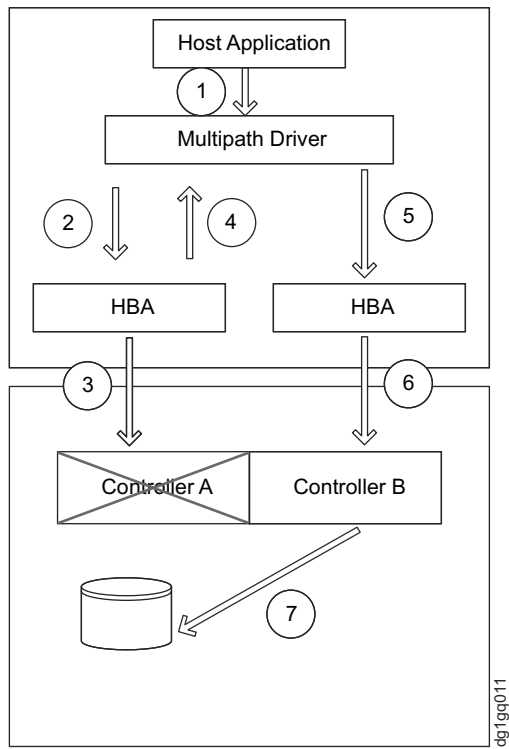


Figure 16. Failover of I/O in a single path environment




Figure 17. Failover of I/O in a multipath environment

Depending on the controller firmware and the multipath driver code, the multipath driver performs different actions for the controller failover depending on the enable failover mode set by selecting the appropriate host as per Table 22 on page 114. There are three controller failover modes depending on the versions of controller firmware:

1. Automatic Volume Transfer (AVT/ADT) failover mode - If the host type is set to enable AVT/ADT failover mode, the multipath driver will redirect I/Os to the surviving controller. You can set the surviving controller to take ownership of the logical drive and process I/Os. Ownership can be set regardless of whether the failed controller is up and running. This is similar to the case in which all paths to a controller fail, or the controller itself fails. Controller firmware version 7.77.xx.xx or earlier supports this failover mode.
2. RDAC failover mode - If the host type is set to **disable AVT/ADT** or **non-ALUA**, the multipath driver will issue a mode page 2C to the surviving controller to move the ownership of the logical drive to the surviving controller. Then, the surviving controller will take the ownership of the logical drive and process I/Os on it. The surviving controller will take ownership of the logical drive no matter whether the other controller is up and running or not as in the cases where all paths to the controller fail or the controller is itself failed. This failover mode is supported with all versions of controller firmware.
3. Asymmetric Logical Unit Access (ALUA) mode - With controller firmware version 7.83.xx.xx and later, if the host type is set to enable ALUA, the multipath driver will just redirect I/Os to the surviving controller. If the "FAILED" controller is up and running as in the case where the paths to the

controller failed but the controller is itself still optimal, the surviving controller will ship the IO to the “FAILED” controller for processing instead of taking the ownership of the logical drive and process I/Os on it. If this condition persists more than 5 minutes, the surviving controller will stop IO shipping to the other controller for processing and take the ownership and the processing of the I/Os to the logical drive.

The advantages of ALUA are:

- “Boot from SAN” server will not failed during boot because the boot LUN is not the path or is not owned by the controller that is on the path that the server first scans during the server boot process. Boot from SAN server is the server whose operating system disk resides in one of the logical drive of the storage subsystem instead of internally inside the server chassis.
- Eliminating unnecessary logical drive failovers/failbacks if there are intermittent shot duration (<5 mins) path interruptions.
- Preventing “LUN ping-pong” in certain conditions where the logical drives are mapped to servers in a cluster environment.
- The logical drive is operated as active-active in a dual controller configuration. The I/Os can be sent to both controllers for processing irrespective of which controller owns the logical drive. In RDAC or AVT/ADT failover mode, only the controller that owns the logical drive can process I/Os to that logical drive. This is also referred to as active-passive operating mode in a dual controller configuration.




Figure 18. All paths to controller fail in AVT/ADT and RDAC failover modes

Figure 18 illustrates the failover when all paths to the controller fail but the controller is itself optimal in AVT/ADT and RDAC failover modes. In this failover scenario, the logical drive ownership is transferred to controller B and controller B processes all I/Os to the logical drives even when controller A is up and optimal

and the failure is caused only by the path failure from the host to controller A. Figure 19 and Figure 20 on page 113 illustrate failover when all paths to the controller fail but the controller is itself optimal in ALUA failover mode.

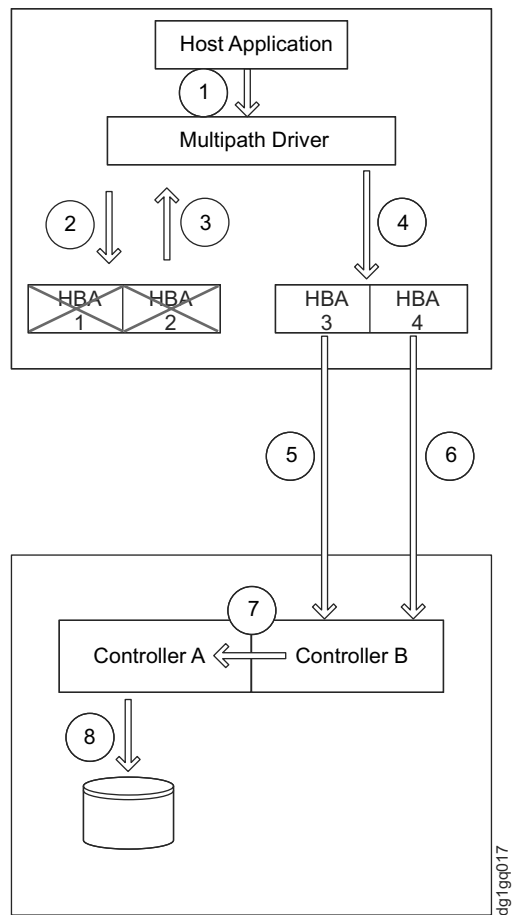


Figure 19. All paths to controller fail in ALUA failover mode. First five minutes of the failover.

During the first five minutes of the failover, the I/Os to the logical drive are shipped internally to Controller A for processing as shown in Figure 19. Controller A is still the owner of the logical drive. After five minutes, if the path to Controller A is still failed, Controller B takes the ownership and the processing of the I/Os to the logical drive as shown in Figure 20 on page 113.




Figure 20. All paths to the controller fail in ALUA mode. Five minutes into the failure

When operating in failover modes 1 and 2 above, the dual controller in the storage subsystem operate in an active-passive combination from the mapped LUN perspective. This means I/O can be only sent to the controller that owns the mapped LUN for processing. The other controller will be in standby mode until either the LUN-owning controller failed or all paths to the LUN-owning controller failed. I/Os sent to the controller that does not own the mapped LUN will either cause the LUN to failover that controller (AVT/ADT mode) or be failed by the controller (RDAC mode). In ALUA failover mode, the dual controllers are now operated as active-active combination from the mapped LUN perspective. I/Os can be sent to both controllers for processing instead of just to the owning controller. The LUN non-owning controller does not have to operate in standby/passive mode until the LUN-owning controller failed. The I/Os are automatically routed internally to the controller that owns the LUNs for processing. In addition, the LUN ownership changes only when one controller processes more than 75% of the I/Os to the LUN within a 5 minute period in ALUA mode.

Controller firmware versions 7.77.xx.xx and earlier support AVT/ADT and RDAC failover modes. Controller firmware version 7.83.xx.xx and later support only RDAC and ALUA failover modes. AVT/ADT mode is not supported in controller firmware version 7.83.xx.xx or later. Note the same controller NVSRAM bit in the host type region is used to enable AVT/ADT or ALUA. Depending on the version of the controller firmware, that bit is either enabled AVT/ADT or ALUA failover mode. To enable which failover mode, the appropriate host type must be selected

for the server host partition. The following table lists the host type for various OS-es and which fail over mode that was enabled for that host type:

Table 22. Failover mode for each Operating System

Host Index	Host type (full name)	Host type (short name) ₁	ADT/AVT ₂	RDAC	ALUA ₂
0	Default	Base	No	Yes	No
1	MacOS	MacOS	No	Yes	No
2	Windows Server 2003/Server 2008 Non-Clustered <Windows>	W2KNETNCL/Windows	No	Yes	No
3	Windows Server 2003/Server 2008 Non-Clustered <Windows>	W2KNETNCL/Windows Clustered	No	Yes	No
4	AIX with Veritas DMP	AIXAVT	Yes	No	No
5	Linux/ Linux with Veritas DMP	LNHAVT/Linux	Yes	No	No
6	AIX	AIX	No	Yes	No
7	HP-UX	HPX	Yes	No	No
8	Solaris (Sparc) <Solaris (with or without MPXIO)>	SOL	No	Yes	No
9	Windows Server 2003/Server 2008 Clustered Non-Clustered (DMP support)	W2KNETNCLDMP <Windows DMP>	Yes	No	No
10	Unused 10/Irix ³	Unused10/Irix	No	Yes	No
11	Unused 11/Netware Failover ³	Unused11/Netware	No	Yes	No
12	IBM TS SAN VCE	IBM TS SAN VCE	Yes	No	No
13	Linux Cluster	LNXCUSTER/ LNXCCLVMWARE <Linux Cluster>	No	Yes	No
14	Solaris (with Veritas DMP)	SOLAVT	Yes	No	No
15	Windows Server 2003/Server 2008 Clustered Clustered (DMP Support)	W2KNETCLDMP	Yes	No	No
16	VMWARE <VMWare>	VMWARE	Yes	No	No
17	HP-UX TPGS	HPXTPGS	No	Yes	No
18	Linux <Linux Non-ADT>	LNK <Linux Non-ADT>	No	Yes	No
19	IBM I/Os	IBM i	No	Yes	No
20	Onstor	Onstor	Yes	No	No

Table 22. Failover mode for each Operating System (continued)

Host Index	Host type (full name)	Host type (short name) ₁	ADT/AVT ₂	RDAC	ALUA ₂
21	Windows ALUA	W2KALUA	No	No	Yes
22	Linux ALUA	LNHALUA	No	No	Yes
23	AIX ALUA w/ TPGS	AIXATPGSLUA	No	No	Yes
24	VMWARE ALUA w/ TPGS	VMWareTPGSALUA	No	No	Yes
25	Mac OS ALUA w/ TPGS	MacOSTPGSALUA	No	No	Yes
26	HP-UX ALUA w/ TPGS	HPXTPGSALUA	No	No	Yes
27	Solaris ALUA w/ TPGS	SOLTPGSALUA	No	No	Yes

Note:

1. The actual name might be slightly different depending on the version of the NVSRAM file loaded. However, the host type index should be the same across all versions.
2. Even though the same NVSRAM bit enables either ADT/AVT depending on the controller firmware, only ALUA the host types (host index 21-27) must be used for enabling ALUA failover mode because of additional ALUA-specific settings are required.
3. Irix and Netware failover host types are defined in the NVSRAM files for controller firmware version 7.77.xx.xx or earlier. For controller firmware 7.83.xx.xx or later, Netware and Irix servers are not supported as host attachment; therefore, these host type were changed to 'Unused'.

Failback

The multipath drivers also monitor the status of the failed paths periodically and failback the logical drive to the preferred controller once the failed path is restored. In the case that some of multiple paths to the controller failed and then restored, the multipath driver will start using the restored path again to send I/Os. The multipath driver use the same mode (AVT/ADT, RDAC or ALUA) as described in the failover section to move the logical drive back to preferred controller.

The automatic logical drive failback feature of the multipath driver could be disabled in server clustered configurations to prevent 'LUN ping-pong' between controllers problem in certain failover scenarios.

Using multipath drivers to automatically manage logical drive fail-over and fail-back

Host systems that are attached to the DS3000, DS4000, or DS5000 storage subsystem for I/O activity require a multipath driver (sometimes referred to as an RDAC or failover driver) for Fibre Channel path redundancy. The multipath driver monitors I/O paths. If a component failure occurs in one of the Fibre Channel paths, the multipath driver reroutes all I/O to a different path. Your multipath driver depends on the operating system that you have installed.

In the Microsoft Windows environment another multipath driver, referred to as Windows RDAC, was previously provided with Storage Manager host software version 9 and earlier. Support for Windows RDAC was terminated with the release of Storage Manager host software version 10, and later in conjunction with controller firmware version 7.xx.xx.xx and later. In addition, support for AIX fcp_array is being phased out. AIX fcp_array users must migrate to the AIX MPIO multipath driver at the earliest time window.

An IBM Fibre Channel host bus adapter (HBA) provides the interface between a host server and a storage subsystem. Storage subsystem Fibre Channel HBAs are high-performance, direct-memory access, bus-master host adapters that are designed for high-end systems. These HBAs support all Fibre Channel peripheral devices that support private-loop, direct-attach, and fabric-loop attachment. The IBM Host Adapter device driver enables your operating system to communicate with the Fibre Channel HBA.

Table 16 lists the multipath driver or drivers that are supported for different operating systems. Refer to the SSIC to determine which multipath driver is supported for a certain OS version for a particular storage subsystem model.

Table 23. Multipath driver by operating system


Operating system	Multipath driver
AIX	fcp_array (also called RDAC), MPIO, or SDDPCM
HP-UX	LVM , native multipathing, or IBM SDD
Linux	MPP (also called Linux RDAC or RDAC), Veritas DMP, or native (in-distro) Linux Device Mapper Multipath (DM-Multipath)
Mac OS	Provided by the HBA manufacturer
NetWare	Novell MPE
Solaris	RDAC, MPxIO or Veritas DMP
SVC	SDD
VMware	NMP
Windows	MPIO DSM or Veritas DMP DSM

With the exception of Windows MPIO, multipath driver files are not included on the Storage Manager DVD. Check the SSIC and the Storage Manager readme file for the minimum file set versions that are required for your operating system. To learn how to find the readme files on the web, see “Finding Storage Manager software, controller firmware, and readme files” on page xiv. To install the multipath driver, follow the instructions in “Installing a multipath driver” on page 120.

Multipathing refers to the ability of the host to recognize multiple paths to the storage device. This is done by using multiple HBA ports or devices within the host server that are connected to SAN fabric switches, which are also connected to the multiple ports on the storage devices. For the storage products that are referred to as DS3000, DS4000, or DS5000, these devices have two controllers within the storage subsystem that manage and control the disk drives. These controllers behave in either active or passive fashion. Ownership and control of a particular LUN is done by one controller. The other controller is in a passive mode until a

failure occurs, at which time the LUN ownership is transferred to that controller. Each controller might have more than one fabric port for connectivity to the SAN fabric.

Figure 21 shows a sample multipath configuration for all supported operating systems except AIX fcp_array and Solaris RDAC multipath configurations. Figure 22 on page 118 shows a sample multipath configuration for the AIX fcp_array, Microsoft Windows RDAC (no longer supported), and Solaris RDAC multipath configurations.



See “Drives supported by IBM System Storage DS Storage Manager” on page 56 for more information.
Figure 21. Host HBA to storage subsystem controller multipath sample configuration for all multipath drivers except AIX fcp_array and Solaris RDAC

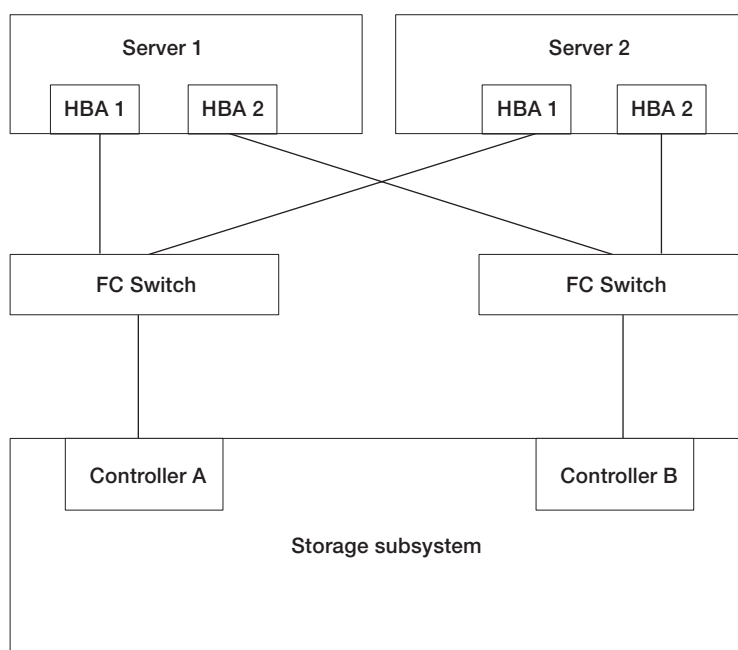


Figure 22. Host HBA to storage subsystem controller multipath sample configuration for AIX fcp_array and Solaris RDAC multipath drivers

Most multipath drivers can support multiple paths. Table 24 shows the number of paths each driver can support. Note that the AIX fcp_array and Solaris RDAC can support only two paths, one to each controller.

Table 24. Number of paths each multipath driver supports by operating system

Driver	Number of paths	Default
AIX MPIO	Unlimited	Not applicable
AIX RDAC	2	Not applicable
HP-UX native	65,536	Not applicable
HP-UX PVlinks	8,192	Not applicable
Linux MPP	Unlimited	4
Linux Veritas DMP	Unlimited	Not applicable
Mac OS	Unlimited	Not applicable
Solaris MPxIO	Unlimited	Not applicable
Solaris RDAC	2	Not applicable
Solaris Veritas DMP	Unlimited	Not applicable
SVC	32	Not applicable
VMware	Unlimited - 8 or fewer for best performance	Not applicable
Windows MPIO DSM	32 paths per LUN, 16 per controller	4
Windows Veritas DMP DSM	Unlimited	Not applicable

Using host bus adapters

This section provides a basic overview of host bus adapters (HBAs), as well as instructions for connecting HBAs in a Fibre Channel switch environment.

Understanding host bus adapters

Host bus adapters (HBAs) are used to connect servers to Fibre Channel topologies. The function of an HBA is similar to that of network adapters used to access LAN resources. The device driver for an HBA is typically responsible for providing support for a Fibre Channel topology, whether point-to-point, loop, or fabric. The DS3000, DS4000, and DS5000 storage subsystems support the Fibre Channel (FC), Serial Attached SCSI (SAS), Fibre Channel over Ethernet (FCoE) and iSCSI host bus adapters (HBAs) to connect the host servers to the storage subsystems. These storage subsystems also support iSCSI via the regular Ethernet NIC adapters. However, all HBAs are not supported on DS3000, DS4000, and DS5000 storage subsystems. For detailed HBA support information, go to <http://www.ibm.com/systems/support/storage/config/ssic>.

See documentation for information about HBA settings that can be customized for a certain operating-system environment. This documentation also includes instructions about how to change these settings. Caution should be made in changing these settings because an incorrect setting can cause degradation in performance or intermittent failures. Also see the readme file that is included in the host bus adapter BIOS, firmware, or device driver package for any up-to-date changes to the settings. The tables in Appendix A, “Host bus adapter settings,” on page 271 show the required values of selected settings for various HBA vendors.

Connecting HBAs in a Fibre Channel switch environment

There are two primary zoning schemes you can use when you connect Fibre Channel host bus adapters (HBAs) in host servers to storage subsystem host ports in a Fibre Channel switch environment. In a one-to-one zoning scheme, each HBA port is zoned to one controller host port. In a one-to-two zoning scheme, each HBA port is zoned to two controller host ports.

As a general rule, the HBA and the storage subsystem host port connections must be zoned to minimize the possible interactions between the ports in a SAN fabric environment. A one-to-one zoning scheme, though not required, minimizes interactions because it connects one HBA port to just one server host port. However, the zoning scheme you choose depends on your host-storage SAN fabric topology and the capabilities of your Fibre Channel switches.

Depending on your host-storage SAN fabric topology Fibre Channel switch capabilities, you can implement one of the two following zoning schemes in Figure 23 on page 120 and Figure 24 on page 120.

Note: For more information about zoning best practices and requirements, see the *Fibre Channel Switch Hardware Reference Guide* or other documentation that came with the Fibre Channel switch. For links to switch documentation on the IBM website, go to

www.ibm.com/servers/storage/support/san/index.html

In this zoning scheme (denoted by the translucent bar), one HBA port is zoned to one controller host port.

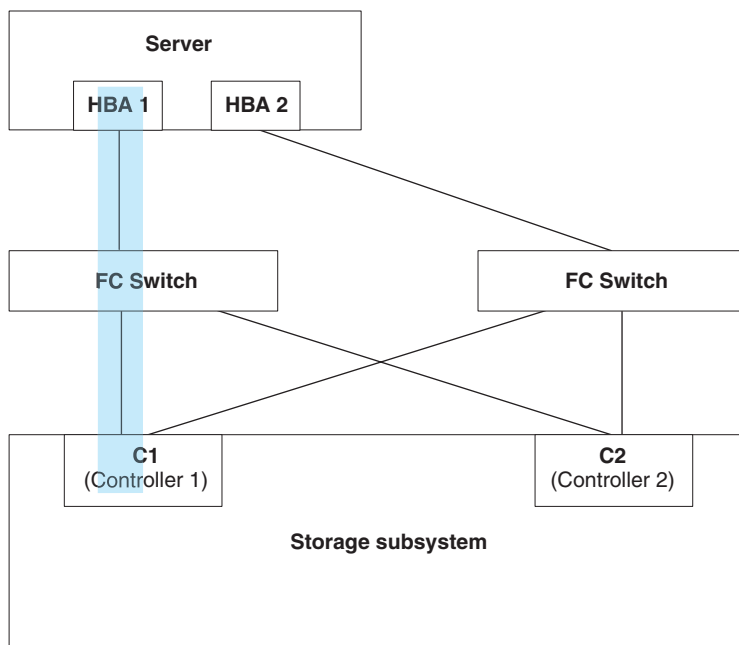


Figure 23. One-to-one zoning scheme

In this zoning scheme (denoted by the translucent bars), one HBA port is zoned to two controller host ports.




Figure 24. One-to-two zoning scheme

Installing a multipath driver

You must install a multipath driver on all hosts that are attached to your storage subsystem, regardless whether these hosts will have multiple paths to the storage subsystem. This section describes various multipath drivers, how to check the current multipath driver program driver version level, how to update the multipath device driver, and how to verify that the multipath update is complete.

Windows MPIO or MPIO/DSM multipath driver Overview

This multipath driver is included in the Storage Manager host software package for Windows. MPIO is a DDK kit from Microsoft for developing code that manages multipath devices. The DDK kit contains a core set of binary drivers that are installed with the storage subsystem Device Specific Module (DSM); the drivers are designed to provide a transparent system architecture that relies on Microsoft Plug and Play. These binary drivers provide LUN multipath functionality and maintain compatibility with existing Microsoft Windows device driver stacks simultaneously. For Windows Server 2003, the MPIO is installed with the MPIO DSM. In Windows Server 2008, only the MPIO DSM is installed because Windows 2008 comes with MPIO installed.

The MPIO driver performs the following tasks:

- Detects and claims the physical disk devices that are presented by the storage subsystems (according to Vendor or 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 (IBM) customized Device-Specific Module (DSM) for the following information about storage subsystem behavior:
 - I/O routing information
 - Conditions that require a request to be retried, failed, failed over, or fail back; for example, Vendor-Unique errors
 - Miscellaneous functions such as Release or Reservation commands

Multiple DSMs for different disk storage subsystems can be installed in the same host server as long as they do not claim the same disk device.

Native SCSI-2 Release/Reservation commands in a multipath environment

If multiple paths exist to a single controller and a SCSI-2 release/reservation (R/R) is received for a logical drive, the MPIO DSM driver selects one path to each controller and repeats the request (called a reservation path). This function is necessary because the controllers cannot accept SCSI-2 R/R requests through multiple paths for a given logical drive. After the reservation path has been established, subsequent I/O requests for logical drives are restricted to that path until a SCSI-2 release command is received. The MPIO DSM driver distributes the reservation paths if multiple logical drives are mapped to the host, which distributes the load across multiple paths to the same controller.

Note: This method of SCSI reservation handling is incompatible with the controller ALUA feature. It must not be used when that feature is enabled.

Translating SCSI-2 Release/Reservation commands to SCSI-3 persistent reservations

The MPIO DSM driver also supports the ability to translate the SCSI-2 R/R commands into SCSI-3 persistent reservations. This function allows a logical drive to use one of the previously mentioned load-balancing policies across all of the available controller paths rather than being restricted to a single reservation path.

This feature requires the MPIO DSM driver to establish a unique “reservation key” for each host. This key is stored in the Registry and is named S2toS3Key. If this key is present, translations are performed, or else the “cloning” method is used.

Per-Protocol I/O timeout values

The timeout value associated with a non-passthrough I/O requests, such as read/write requests, is based on the Microsoft disk driver's **TimeOutValue** parameter, as defined in the Registry. A feature within the DSM allows a customized timeout value to be applied based on the protocol (such as Fibre Channel, SAS, or iSCSI) that a path uses. Per-protocol timeout values provide these benefits:

- Without per-protocol timeout values, the **TimeOutValue** setting is global and affects all storage.
- The **TimeOutValue** is typically reset when an HBA driver is upgraded.
- For Windows Server 2003, the default disk timeout value can be adjusted based on the size of the I/O request. Adjusting the default disk timeout value helps support legacy SCSI devices.
- The DSM customized timeout feature allows a more predictable timeout setting for Windows Server 2003 environments. For information about the configurable parameters for the customized timeout feature, go to Configuration Settings for the Windows DSM and the Linux RDAC.

The per-protocol timeout values feature slightly modifies the way in which the **SynchTimeout** parameter is evaluated. The **SynchTimeout** parameter determines the I/O timeout for synchronous requests generated by the DSM driver.

Examples include the SCSI-2 to SCSI-3 PR translations and inquiry commands used during device discovery. It is important that the timeout value for the requests from the DSM driver be at least as large as the per-protocol I/O timeout value. When a host boots, the DSM driver performs these actions:

- If the value of the **SynchTimeout** parameter is defined in the Registry key of the DSM driver, record the current value.
- If the value of the **TimeOutValue** parameter of the Microsoft disk driver is defined in the Registry, record the current value.
- Use the higher of the two values as the initial value of the **SynchTimeout** parameter.
- If neither value is defined, use a default value of 10 seconds.
- For each synchronous I/O request, the higher value of either the per-protocol I/O timeout or the **SynchTimeout** parameter is used. For example:
 - If the value of the **SynchTimeout** parameter is 120 seconds, and the value of the **TimeOutValue** parameter is 60 seconds, 120 seconds is used for the initial value.
 - If the value of the **SynchTimeout** parameter is 120 seconds, and the value of the **TimeOutValue** parameter is 180 seconds, 180 seconds is used for the initial value of the synchronous I/O requests for the DSM driver.
 - If the I/O timeout value for a different protocol (for example, SAS) is 60 seconds and the initial value is 120 seconds, the I/O will be sent using a 120-second timeout.

Selective LUN transfer

This feature limits the conditions under which the DSM driver will move a LUN to the alternative controller to three cases:

1. When a DSM driver with a path to only one controller, the non-preferred path, discovers a path to the alternate controller.
2. When an I/O request is directed to a LUN that is owned by the preferred path, but the DSM driver is attached to only the non-preferred path.
3. When an I/O request is directed to a LUN that is owned by the non-preferred path, but the DSM driver is attached to only the preferred path.

Case 2 and case 3 have these user-configurable parameters that can be set to tune the behavior of this feature.

- The maximum number of times that the LUN transfer will be issued. This parameter setting prevents a continual ownership thrashing condition from occurring in cases where the controller enclosure or the controller-drive enclosure is attached to another host that requires the LUN be owned by the current controller.
- A time delay before LUN transfers are attempted. This parameter is used to de-bounce intermittent I/O path link errors. During the time delay, I/O requests will be retried on the current controller to take advantage of the possibility that another host might transition the LUN to the current controller.

For further information about these two parameters, go to Configuration Settings for the Windows DSM and the Linux RDAC.

In the case where the host system is connected to both controllers and an I/O is returned with a 94/01 status, in which the LUN is not owned and can be owned, the DSM driver modifies its internal data on which controller to use for that LUN and reissue the command to the other controller. To avoid interfering with other hosts that might be attached to a controller enclosure or the controller-drive enclosure, the DSM driver does not issue a LUN transfer command to that controller enclosure or controller-drive enclosure.

When the DSM detects that a logical drive transfer operation is required, the DSM will not immediately issue the failover/failback command. It will, with the default settings, delay for three seconds before sending the command to the storage subsystem. This delay provides time to batch together as many logical drive transfer operations for other LUNs as possible. The controller single-threads logical drive transfer operations and will therefore reject additional transfer commands until the controller has completed the operation that it is currently working on. This results in a period during which I/Os are not successfully serviced by the storage subsystem. By introducing a delay during which logical drive transfer operations can be aggregated into a batch operation, the DSM reduces the likelihood that a logical drive transfer operation will exceed the retry limit. In large system configurations, you might need to increase the default three-second delay value because, with more hosts in the configuration, more logical drive transfer commands might be sent.

The delay value is located at: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\<DSM_Driver>\Parameters\LunFailoverDelay.

This feature is enabled when these conditions exist:

- The controller enclosure or the controller-drive enclosure does not have AVT enabled.
- The DSM driver configurable parameter **ClassicModeFailover** is set to 1.
- The DSM driver configurable parameter **DisableLunRebalance** is set to 4.

Windows failover cluster

Clustering for the Windows Server 2008 OS and the Windows Server 2008 R2 OS uses SCSI-3 persistent reservations natively. As a result, the DSM driver does not perform translations for any SCSI-2 R/R commands, and you can use one of the previously mentioned load-balancing policies across all controller paths. Translations still occur if the DSM driver is running in a Windows Server 2003 OS-based environment. If you are operating in a clustered environment and not utilizing the I/O Shipping feature of the CFW or the Selective LUN Transfer feature of the DSM, set the **DisableLunRebalance** parameter to 3. For information about this parameter, go to Configuration Settings for the Windows DSM and the Linux RDAC.

I/O shipping feature for Asymmetric Logical Unit Access (ALUA)

The I/O Shipping feature implements support for ALUA. With earlier releases of the controller firmware (CFW), the device specific module (DSM) had to send input/output (I/O) requests for a particular logical drive to the controller that owned that logical drive. A controller would reject requests it received for a logical drive that it did not own. This behavior was necessary in order for the storage subsystem to maintain data consistency within the logical drive. This same behavior, however, was responsible for several areas of contention during system boot and during multi-host\path failure conditions.

With the I/O Shipping feature, a storage subsystem can service I/O requests through either controller in a duplex configuration. There is a performance penalty when the non-owning controller accesses a logical drive. In order to maintain the best I/O subsystem performance, the DSM interacts with the CFW to insure that I/O requests are sent to the owning controller if that controller is available.

When you install or update the DSM, by default, the Selective LUN Transfer (SLT) feature is enabled to support I/O Shipping. Some registry values are modified during the DSM update if the previous version did not enable SLT. To prevent enabling SLT so that your storage subsystem will operate without the I/O Shipping feature, edit the registry to have the following settings:

- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\<DSM_Driver>\Parameters\DisableLunRebalance where <DSM_Driver> is the name of the DSM driver used in your storage subsystem.
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\<DSM_Driver>\Parameters\ClassicModeFailover where <DSM_Driver> is the name of the DSM driver used in your storage subsystem.

Storport Miniport HBA device driver

For Windows operating systems, the Storage Manager provides the MPIO DSM device driver that is based on the Microsoft Storport Miniport device driver model.

The Storport Miniport device driver model was introduced in the Microsoft Windows Server 2003 release as a replacement for the SCSIport Miniport device driver model. It is the only supported device driver model for Windows Server 2003 and Windows Server 2008 editions, supporting the AMD64 and EM64T

servers. It does not support the `buschange=0` parameter to bypass the Microsoft Windows operating system Plug and Play driver. Instead, it works with the Plug and Play driver to detect the removal and insertion of devices at the Fibre Channel host bus adapter port.

Only the DS4100, DS4200, DS4300 (standard or turbo models), DS4400, DS4500, DS4700, and DS4800 storage subsystems support this Storport-based device driver. The DS4100, DS4300 (standard or turbo models), DS4400, and DS4500 storage subsystem models must have controller firmware version 6.12.27.xx or later.

See the Storage Manager readme file for Microsoft Windows operating systems for any other additional requirements, such as controller firmware versions or updates.

SCSIport Miniport HBA device driver

For the Windows 2000 operating-system environment, only the device drivers based on the SCSIport Miniport device driver (not the Storport model) are supported.

In previous SCSIport device driver releases, the `buschange=0` parameter allows the RDAC multipath driver to control and monitor device insertion and removal from the HBA port by preventing the Microsoft Plug and Play device driver from managing the HBA port. This new SCSIport device driver version that is used with MPIO does not support the `buschange=0` parameter.

Important: Not all DS4000 and DS5000 controller firmware versions support this functionality. Only DS4000 and DS5000 controller firmware versions 06.12.27.xx (and later) for DS4300 standard or turbo models, and DS4500 storage subsystems or versions 6.16.8x.xx (and later) for DS4200, DS4700, and DS4800 storage subsystems support the SCSIport Miniport device driver.

Before you install the device driver, check the readme file that is included in the device driver package file, as well as the readme file included with the Storage Manager host software for Windows, to see which device drivers and controller firmware versions are supported for DS3000, DS4000, or DS5000 storage subsystems. See “Finding Storage Manager software, controller firmware, and readme files” on page xiv to learn how to access the most recent Storage Manager readme files on the web. Follow the readme device driver installation instructions that are associated with your operating system.

Note: Read the device driver readme file for any required modifications to the default HBA BIOS and host operating system registry settings to provide optimal performance. If you make any changes to the HBA BIOS settings, the server must be rebooted for the changes to be enabled.

For more information, see the documentation that came with your Fibre Channel HBA.

Using dsmUtil

The dsmUtil utility is a command-line driven utility that works only with the Multipath I/O (MPIO) Device Specific Module (DSM) solution. The utility is used primarily as a way to instruct the DSM driver to perform various maintenance tasks, but the utility can also serve as a troubleshooting tool when necessary.

To use the dsmUtil utility, type this command, and press Enter:

```

dsmUtil [[-a [target_id]]
[-c array_name | missing]
[-d debug_level] [-e error_level] [-g
virtual_target_id]
[-o [[feature_action_name=value]] |
[feature_variable_name=value]][, SaveSettings]] [-M]
[-P [GetMpioParameters | MpioParameter=value | ...]]
[-R]
[-s "failback" | "avt" | "busscan" | "forcerebalance"]
[-w target_wwn, controller_index]

```

Type `dsmUtil` without any parameters shows the usage information. Refer to the following table for the `dsmUtil` parameters.

Table 25. *dsmUtil* parameters

Parameter	Description
-a [<i>target_id</i>]	Shows a summary of all storage subsystems seen by the DSM. The summary shows the <i>target_id</i> , the storage subsystem WWID, and the storage subsystem name. If <i>target_id</i> is specified, DSM point-in-time state information appears for the storage subsystem. On UNIX operating systems, the virtual HBA specifies unique target IDs for each storage subsystem. The Windows MPIIO virtual HBA driver does not use target IDs. The parameter for this option can be viewed as an offset into the DSM information structures, with each offset representing a different storage subsystem. For use by Customer and Technical Support representatives only.
-c <i>array_name</i> missing	Clears the WWN file entries. This file is located in Program Files\DSMDrivers\mppdsm\WWN_FILES with the extension .wwn. If the <i>subsystem_name</i> keyword is specified, the WWN file for the specific storage subsystem is deleted. If the <i>missing</i> keyword is used, all WWN files for previously attached storage subsystems are deleted. If neither keyword is used, all of the WWN files, for both currently attached and previously attached storage subsystems, are deleted.

Table 25. *dsmUtil* parameters (continued)

Parameter	Description
<code>-d debug_level</code>	<p>Sets the current debug reporting level. This option only works if the RDAC driver has been compiled with debugging enabled. Debug reporting is comprised of two segments. The first segment refers to a specific area of functionality, and the second segment refers to the level of reporting within that area. The <code>debug_level</code> is one of these hexadecimal numbers:</p> <ul style="list-style-type: none"> • 0x20000000 – Shows messages from the RDAC driver’s initialization routine. • 0x10000000 – Shows messages from the RDAC driver’s device discovery routine. • 0x08000000 – Shows messages from the RDAC driver’s <code>ioctl()</code> routine. 0x04000000 – Shows messages from the RDAC driver’s device open routine (Linux platforms only). • 0x02000000 – Shows messages from the RDAC driver’s device read routine (Linux platforms only). • 0x01000000 – Shows messages related to HBA commands. • 0x00800000 – Shows messages related to aborted commands. • 0x00400000 – Shows messages related to panic dumps. • 0x00200000 – Shows messages related to synchronous I/O activity. • 0x00100000 – Shows messages related to failover activity. • 0x00080000 – Shows messages related to failback activity. • 0x00040000 – Shows additional messages related to failback activity. • 0x00010000 – Shows messages related to device removals. • 0x00001000 – Shows messages related to SCSI reservation activity • 0x00000400 – Shows messages related to path validation activity. • 0x00000001 – Debug level 1. • 0x00000002 – Debug level 2. • 0x00000004 – Debug level 3. • 0x00000008 – Debug level 4. <p>You can combine these options with the logical or operator to provide multiple areas and levels of reporting as needed. For use by Customer and Technical Support representatives only.</p>

Table 25. *dsmUtil* parameters (continued)

Parameter	Description
<code>-e error_level</code>	<p>Sets the current error reporting level to <i>error_level</i>, which can have one of these values:</p> <ul style="list-style-type: none"> • 0 – Show all errors. • 1 – Show path failover, controller failover, retryable, fatal, and recovered errors. • 2 – Show path failover, controller failover, retryable, and fatal errors. • 3 – Show path failover, controller failover, and fatal errors. This is the default setting. • 4 – Show controller failover and fatal errors. • 5 – Show fatal errors. <p>For use by Customer and Technical Support representatives only.</p>
<code>-g target_id</code>	<p>Displays detailed information about the state of each controller, path, and LUNs for the specified storage subsystem. You can find the <i>target_id</i> by running the <code>dsmUtil -a</code> command.</p>
<code>-M</code>	<p>Displays the MPIIO disk-to-drive mappings for the DSM. The output is similar to that found with the <code>SMdevices</code> utility. For use by Customer and Technical Support representatives only.</p>
<code>-o [[feature_action_name[=value]] [feature_variable_name=value]][, SaveSettings]</code>	<p>Troubleshoots a feature or changes a configuration setting. Without the <code>SaveSettings</code> keyword, the changes only affect the in-memory state of the variable. The <code>SaveSettings</code> keyword changes both the in-memory state and the persistent state. Some example commands are:</p> <ul style="list-style-type: none"> • <code>dsmUtil -o</code> – Displays all the available feature action names. • <code>dsmUtil -o DisableLunRebalance=0x3</code> – Turns off the DSM-initiated storage subsystem LUN rebalance (affects only the in-memory state).
<code>-P [GetMpioParameters MpioParameter=value ...]</code>	<p>Displays and sets MPIIO parameters. For use by Customer and Technical Support representatives only.</p>
<code>-R</code>	<p>Remove the load-balancing policy settings for inactive devices.</p>

Table 25. *dsmUtil* parameters (continued)

Parameter	Description
-s ["failback" "avt" "busscan" "forcerebalance"]	Manually initiates one of the DSM driver's scan tasks. A "failback" scan causes the DSM driver to reattempt communications with any failed controllers. An "avt" scan causes the DSM driver to check whether AVT has been enabled or disabled for an entire storage subsystem. A "busscan" scan causes the DSM driver to go through its unconfigured devices list to see if any of them have become configured. A "forcerebalance" scan causes the DSM driver to move storage subsystem logical drives to their preferred controller and ignores the value of the DisableLunRebalance configuration parameter of the DSM driver.
-w target_wwn, controller_index	For use by Customer and Technical Support representatives only.

Veritas DMP DSM driver

See the Symantec Storage Foundation for Windows documentation for instructions for installing the Veritas DMP DSM driver at <http://www.symantec.com/business/support/>.

AIX multipath drivers

An AIX host system requires either the AIX Redundant Disk Array Controller (RDAC) or the MPIO failover driver for Fibre Channel path redundancy. In supported Veritas environments, RDAC is the supported failover driver.

The failover driver monitors I/O paths. If a component failure occurs in one of the Fibre Channel paths, the failover driver reroutes all I/O to another path.

Note: AIX supports both Redundant Disk Array Controller (RDAC) and Multiple Path I/O. These multipath drivers are part of the native AIX operating system. See the AIX documentation for details about the installation of these drivers.

Linux Device Mapper Multipath driver

The Device Mapper Multipath (DMMP or DM-MP) is a generic framework for block devices provided by the Linux operating system. It supports concatenation, striping, snapshots, mirroring, and multipathing. The multipath function is a combination of the kernel modules and user space tools.

Important: The host on which the Linux Device Mapper multipath driver is installed should be either Linux non-AVT/non-ADT (Linux) or LinuxCluster (LNXCLUSTER). You can verify the ADT or AVT status in the host type information in the storage subsystem profile. The ADT or AVT status of the host, on which the Linux Device Mapper multipath driver is installed, should be disabled.

The device mapper multipath driver:

- Provides a single block device node for a multipathed logical unit
- Ensures that I/O is re-routed to available paths in case of a path failure
- Ensures that the failed paths are revalidated as soon as possible

- Configures multipaths to maximize performance
- Reconfigures multipaths automatically when events occur
- Provides DMMP features support to a newly added logical unit
- Provides device name persistence for DMMP devices under /dev/mapper/
- Configures multipaths automatically at an early stage of rebooting so that the OS can be installed and rebooted on a multipathed logical unit

Device mapper multipath (DMMP or DM-MP) is supported on SLES11, SLES11 SP1, RHEL 6.0, RHEL 6.1 or their later versions.

Installing the Device Mapper MultiPath driver

Device mapper multipath (DMMP or DM-MP) is supported on SLES11, SLES11 SP1, RHEL 6.0, RHEL 6.1 or their later versions.

Refer to <http://www.ibm.com/systems/support/storage/config/ssic> for information on the DS3000/DS5000 subsystems and controller firmware versions that are supported with DMMP.

Installing the Device Mapper MultiPath on SLES11 base

Note: ALUA functionality is not supported with SLES 11 base. You must upgrade to SLES 11 SP1 or later with controller firmware version 7.8x.xx.xx or later. SLES 11 base operating system does not have all the packages to support device mapper for IBM DS Storage subsystems. It is recommended that you use the most recent versions of the following components, if available. Else, you must have at least the following component versions:

Table 26. Minimum version required for each component

Component	Minimum version	Download location
Kernel	kernel-default-2.6.27.29-0.1.1	http://download.novell.com/patchfinder
scsi_dh_rdac driver	lsi-scsi_dh_rdac-kmp-default-0.0_2.6.27.19_5-1	http://drivers.suse.com/driver-process/pub/update/LSI/sle11/common/x86_64/
Device Mapper library	device-mapper-1.02.27-8.6	http://download.novell.com/patchfinder
Kpartx	kpartx-0.4.8-40.6.1	http://download.novell.com/patchfinder
Multipath_tools	multipath-tools-0.4.8-40.6.1	http://download.novell.com/patchfinder

Ensure that you install all the dependent packages before continuing further. For more details, refer to the SUSE Linux Enterprise Server 11 Installation and Administration Guide in Novel/SuSe website.

Complete the following steps to install device mapper multipath on SLES11 base:

1. Use the media supplied with the operating system vendor to complete the installation of SLES 11.
2. Download and install the errata kernel 2.6.27.29-0.1.

3. Reboot to the 2.6.27.29-0.1 kernel.
4. Install device-mapper-1.02.27-8.6.
5. Install kpartx-tools-0.4.8-40.6.1.
6. Install multipath-tools-0.4.8-40.6.1.
7. Update and configure /etc/multipath.conf. A sample file is stored at /usr/share/doc/packages/multipath-tools/multipath.conf.synthetic. Copy and rename this file to /etc/multipath.conf. Refer to “Working with Multipath.conf file” on page 133 for more details.
8. Enable multipathd service using the following command: **#chkconfig multipathd on.**
9. Edit the /etc/sysconfig/kernel file to add scsi_dh_rdac to the INITRD_MODULES list. This should add the scsi_dh_rdac to initrd.
10. Install lsi-scsi_dh_rdac-kmp-default-0.0_2.6.27.19_5-1.
11. Reboot the host.

Installing the Device Mapper MultiPath on SLES11 sp1 or later

All the components required for DMMP are included in SUSE Linux Enterprise Server (SLES) version 11.1 sp1 installation media. By default, DMMP is disabled in SLES. Complete the following steps to enable DMMP components on the host:

Note: Asymmetric Logical Unit Access (ALUA) is supported with controller firmware 7.83.xx.xx or later. SLES sp2 and later have the patches required for ALUA as a part of kernel distribution.

1. Use the media supplied with the operating system vendor to complete installation of SLES11 sp1.
2. If you have installed SLES 11 sp2 or later, skip this step. Else, Run `perform rpm -qa | grep <name of the package>` to verify whether the following packages have been installed.

- kpartx-0.4.8-40.21.1.1.09.00.0000.0006
- multipath-tools-0.4.8-40.21.1.1.09.00.0000.0000
- scsi_dh_rdac-kmp-default-09.00.0000.0006_2.6.32.12_0.7-sles11.1

and to enable the ALUA feature, install the packages using the following rpm commands -

- **rpm -ivh kpartx-0.4.8-40.21.1.1.00.00.0000.0005.<arch>.rpm --force**
- **rpm -ivh multipath-tools-0.4.8-40.21.1.1.00.00.0000.0005.<arch>.rpm --force**
- **rpm -ivh scsi_dh_rdac-kmp-default-00.00.0000.000<X>_2.6.32.12_0.7-sles11.1.<arch>.rpm --force**

where <arch> is replaced with the appropriate architecture (x86,x86_x64 , orPPC64)

3. Update and configure /etc/multipath.conf. A sample file is stored at /usr/share/doc/packages/multipath-tools/multipath.conf.synthetic. Copy and rename this file to /etc/multipath.conf. Refer to “Working with Multipath.conf file” on page 133 for more details.
4. Enable multipathd service using the following command: **#chkconfig multipathd on.**
5. Edit the file /etc/sysconfig/kernel to add scsi_dh_rdac to the INITRD_MODULES list. This should add scsi_dh_rdac to initrd.

6. Create the new initrd image using the following command: `# mkinitrd -k /boot/vmlinux-$(uname -r) -i /boot/initrd-$(uname -r)-scsi_dh -M /boot/System.map-$(uname -r)`
7. Update the boot loader configuration file (grub.conf , lilo.conf, or yaboot.conf) with the newly built initrd.
8. Reboot the host to boot with the new initrd image.

Installing the Device Mapper Multi-Path on RHEL 6.0, RHEL 6.1 or later

All the components required for DMMP are included in RHEL 6 and 6.1 installation media. By default, DMMP is disabled. Complete the following steps to enable DMMP components on the host.

Note: Asymmetric Logical Unit Access (ALUA) is not supported on RHEL 6.0. You must upgrade to RHEL 6.1 or later, and the controller firmware must be 7.83.xx.xx or later.

1. Use the media supplied with the operating system vendor to complete installation of RHEL 6.0, RHEL 6.1 or later.
2. If you have installed RHEL 6 update 2, skip this step as RHEL 6 update 2 or later already has the patches required for ALUA support. These patches are a part of kernel distribution. Else, Run `perform rpm -qa | grep <name of the package>` to verify whether the following packages have been installed.

- `kpartx-0.4.8-40.21.1.1.09.00.0000.0006`
- `multipath-tools-0.4.8-40.21.1.1.09.00.0000.0000`
- `scsi_dh_rdac-kmp-default-09.00.0000.0006_2.6.32.12_0.7-sles11.1`

and to enable the ALUA feature, install the packages using the following rpm commands -

- `rpm -ivh kpartx-0.4.8-40.21.1.1.00.00.0000.0005.<arch>.rpm --force`
- `rpm -ivh device-mapper-multipath-libs-0.4.9-41.1.el6.00.00.0000.0005.<arch>.rpm --force`
- `rpm -ivh device-mapper-multipath-0.4.9-41.1.el6.00.00.0000.000<X>.<arch>.rpm --force`
- `rpm -ivh scsi_dh_rdac-kmod-00.00.0000.0005-el6.<arch>.rpm`

where <arch> is replaced with the appropriate architecture (x86,x86_x64 , orPPC64)

3. Update and configure /etc/multipath.conf. A sample file is stored at /usr/share/doc/packages/multipath-tools/multipath.conf.synthetic. Copy and rename this file to /etc/multipath.conf. Refer to “Working with Multipath.conf file” on page 133 for more details.
4. Enable multipathd service using the following command: `#chkconfig multipathd on`
5. Create an initramfs image using the scsi_dh_rdac driver:
 - a. Create a file `scsi_dh_alua.conf` in the `/etc/modprobe.d/` directory.
 - b. In this file, add the following: `alias scsi_hostadapter99 scsi_dh_rdac`
6. Run the following command to create an initramfs image: `#dracut -f /boot/initrd-$(uname -r)-scsi_dh $(uname -r)`
7. Update the boot loader configuration file (grub.conf , lilo.conf, or yaboot.conf) using initramfs.
8. Reboot the host to boot with the new initramfs image.

Working with Multipath.conf file

Multipath.conf is the configuration file for the multipath daemon, multipathd. This file overrides the built-in configuration table for multipathd. Any line in the file that begins with the first non-white-space character # is a comment line. Empty lines should be ignored.

By default, DMMP is supported on certain machine-type models of the IBM DS3000/DS5000 subsystems. However, IBM recommends overriding the default settings using the multipath.conf file to ensure that the DMMP settings are as under:

To set up the multipath.conf file, complete the following steps:

1. Copy the sample multipath.conf from the appropriate directory, depending on whether the Linux operating system is Redhat RHEL or Novell SLES, to the /etc directory.
 - For SLES, the file is named multipath.conf.synthetic and is stored in this directory: /usr/share/doc/packages/multipath-tools/.
 - For RHEL, the file is called multipath.conf.defaults and is stored in this directory: /usr/share/doc/device-mapper-multipath-0.4.9/.
2. Rename the file multipath.conf.
3. Make the configuration changes described in this section to the new /etc/multipath.conf file. The content of the sample multipath.conf file varies, depending on whether it is from SLES or RHEL kernels.

Note: All entries for multipath devices are commented out initially. To uncomment, remove the first character (#) from that section. You must uncomment the three sections - default, blacklist, and devices. The configuration file is divided into five sections:

defaults

Specifies all of the default values.

blacklist

Blacklists new installations. The default blacklist is listed in the commented-out section of the /etc/multipath.conf file. Blacklist the device mapper multipath by WWID if you do not want to use this functionality.

blacklist_exceptions

Specifies any exceptions to the items in the blacklist section.

devices

Lists all of the multipath devices with their matching vendor and product values.

multipaths

Lists all of the multipath devices with their matching WWID values.

To determine the attributes of a multipath device, check the multipaths section of the /etc/multipath.conf file; then the devices section; and then the defaults section. Depending on the version of the Linux kernel, the devices section of the sample multipath.conf file might already have settings defined for your storage subsystem model product ID. All you need to do is verify that the settings match the recommended settings listed below. Otherwise, you have to manually enter the devices settings for your subsystem model product ID. If you have multiple storage subsystems with different product IDs connected to the Linux host, add the device settings for each storage subsystem product ID in the devices section of the

/etc/multipath.conf file. Sample settings for DS3500 (product ID 1746) and DS5100/DS5300 (product ID 1818) in the devices section of the multipath.conf file in SLES operating systems are shown below:

Note: If the Product ID exceeds four characters, use only the first four characters. In the following example, although the Product ID is '1746 FAStT', the product is specified as '1746'. Similarly, '1818 FAStT' is specified as '1818'.

```
Devices {
  device {

        vendor                "IBM"
        product                "1746"
        path_grouping_policy   group_by_prio
        getuid_callout         "/lib/udev/scsi_id -g -u -d /dev/%n"
        path_selector          "round-robin 0"
        path_checker            rdac
        features                "2 pg_init_retries 50"
        hardware_handler       "1 rdac"
        prio                    rdac
        failback                immediate
        no_path_retry           15
        rr_min_io               100
        rr_weight                priorities
    }

  device {

        vendor                "IBM"
        product                "1818"
        path_grouping_policy   group_by_prio
        getuid_callout         "/lib/udev/scsi_id -g -u -d /dev/%n"
        path_selector          "round-robin 0"
        path_checker            rdac
        features                "2 pg_init_retries 50"
        hardware_handler       "1 rdac"
        prio                    rdac
        failback                immediate
        no_path_retry           15
        rr_min_io               100
        rr_weight                priorities
    }
}
```

Sample settings for DS3500 (product ID 1746) and DS5100/DS5300 (product ID 1818) in the devices section of the multipath.conf file in RHEL operating systems are shown below:

```
Devices {
  device {

        vendor                "IBM"
        product                "1746"
        path_grouping_policy   group_by_prio
        getuid_callout         "/lib/udev/scsi_id
                                --whitelisted --device=/dev/%n"
        path_selector          "round-robin 0"
        path_checker            rdac
        features                "2 pg_init_retries 50"
        hardware_handler       "1 rdac"
        prio                    rdac
        failback                immediate
        no_path_retry           15
        rr_min_io               100
        rr_weight                priorities
    }
}
```

```

device {
    vendor                "IBM"
    product               "1818"
    path_grouping_policy  group_by_prio
    getuid_callout        "/lib/udev/scsi_id
                        --whitelisted --device=/dev/%n"
    path_selector         "round-robin 0"
    path_checker          rdac
    features              "2 pg_init_retries 50"
    hardware_handler      "1 rdac"
    prio                  rdac
    failback              immediate
    no_path_retry         15
    rr_min_io             100
    rr_weight             priorities
}

```

If you have Access LUN (sometimes referred to as UTM LUN) mapped to the host partitions, include an entry in the blacklist section of the `/etc/multipath.conf` file, so that the file is not managed by the DMMP. The Storage manager host software uses the Access LUN for in-band management of the storage subsystem. The entries should follow the pattern of the following example:

```

blacklist {
    device {
        vendor "*"
        product "Universal Xport"
    }
}

```

The following table describes the attributes and values in the devices section of the `/etc/multipath.conf` file.

Table 27. Attributes and parameter values in the `multipath.conf` file

Attribute	Parameter value	Description
path_grouping_policy	group_by_prio	This attribute determines the path grouping policy to be applied to this specific vendor and product storage.
prio	rdac	This attribute sets the program and arguments to determine the path priority routine. The specified routine should return a numeric value specifying the relative priority of this path. Higher numbers have a higher priority.
getuid_callout	For SLES"/lib/udev/scsi_id -g -u -d /dev/%n" For RHEL"/lib/udev/scsi_id --whitelisted--device=/dev/%n"	This attribute determines the program and arguments to call out and obtain a unique path identifier.
polling_interval	5	This attribute determines the interval between two path checks, in seconds.
path_checker	rdac	This attribute establishes the method used to determine the state of the path.

Table 27. Attributes and parameter values in the *multipath.conf* file (continued)

Attribute	Parameter value	Description
path_selector	"round-robin 0"	This attribute determines the path selector algorithm to use when there is more than one path in a path group.
hardware_handler	"1 rdac"	This attribute determines the hardware handler to use for handling device-specific knowledge.
failback	immediate	This attribute determines how the daemon manages path group failback. In this example, the parameter is set to 10 seconds, so failback occurs 10 seconds after a device comes online. To disable the failback, set this parameter to manual. Set it to immediate to force failback to occur immediately.
features	"2 pg_init_retries 50"	This attribute enables features. In this example, the kernel parameter <code>pg_init_retries</code> is set to 50. The parameter <code>pg_init_retries</code> is used to retry the mode select commands.
no_path_retry	30	This attribute determines the number of retries before queuing is disabled. Set this parameter to fail for immediate failure (no queuing). When this parameter is set to queue, queuing continues indefinitely.
rr_min_io	100	The number of IO to route to a path before switching to the next in the same path group.
rr_weight	priorities	If set to priorities the multipath configurator will assign path weights as "path prio * rr_min_io"

Using the Device Mapper Devices

Multipath devices are created under the `/dev/` directory with the prefix `dm-`. These devices are the same as any other block devices on the host. To list all of the multipath devices, run the **multipath -ll** command. The following example shows system output from the **multipath -ll** command for one of the multipath devices in non-ALUA failover.

```

mpathp (3600a0b80005ab177000017544a8d6b92) dm-0 IBM, 1746 FASTT
[size=5.0G][features=3 queue_if_no_path pg_init_retries
50][hwandler=1 rdac][rw]
\_ round-robin 0 [prio=6][active]
\_ 5:0:0:0 sdc 8:32 [active][ready]
\_ round-robin 0 [prio=1][enabled]
\_ 4:0:0:0 sdb 8:16 [active][ghost]

```

The following example shows system output from the **multipath -ll** command for one of the multipath devices in ALUA failover.

```

mpathf (3600a0b800047516e00006d864f70696c) dm-11 IBM, 1746 FASTT
size=1.0G features='3 queue_if_no_path pg_init_retries 50' hwandler='1 rdac' wp=rw
| +- policy='round-robin 0' prio=14 status=active
|   ~ 9:0:0:3 sdac 65:192 active ready running
| +- policy='round-robin 0' prio=9 status=enabled
|   ~ 10:0:0:3 sds 65:32 active ready running

```

It will show two active ready paths, with priorities 14 and 9 (if the LUN is on the preferred controller). If the alternate, non-preferred, controller owns the LUN, the priorities will show as 12 and 11, like this:

```

mpathe (3600a0b800029e8320000623d4f70486a) dm-15 IBM, 1746 FASTT
size=1.0G features='3 queue_if_no_path pg_init_retries 50' hwandler='1 rdac' wp=rw
| +- policy='round-robin 0' prio=12 status=active
|   ~ 16:0:0:6 sdah 66:16 active ready running
| +- policy='round-robin 0' prio=11 status=enabled
|   ~ 15:0:0:6 sdat 66:208 active ready running

```

In the preceding example, the multipath device node for this device is `/dev/mapper/mpathp` and `/dev/dm-0`. The following table lists some basic options and parameters for the **multipath** command.

*Table 28. Options and parameters for the **multipath** command*

Command	Description
<code>multipath -h</code>	Print usage information
<code>multipath -ll</code>	Show the current multipath topology from all available information (sysfs, the device mapper, path checkers, and so on)
<code>multipath -f map</code>	Flush the multipath device map specified by the <i>map</i> option, if the map is unused
<code>multipath -F</code>	Flush all unused multipath device maps

Troubleshooting the Device Mapper

Use the information in the following table to troubleshoot the Device Mapper.

Table 29. Troubleshooting the Device Mapper

Situation	Action
Check whether the multipath daemon, <code>multipathd</code> , is running	At the command prompt, run the command <code>/etc/init.d/multipathd status</code>
Determine why no devices are listed when you run the multipath -ll command	At the command prompt, run the command <code>#cat /proc/scsi/scsi</code> . The system output displays all of the devices that are already discovered. Verify that the <code>multipath.conf</code> file has been updated with the proper settings.

Known Issues and Limitations

- When storage is configured with ADT/AVT mode, delays in device discovery might occur. Delays in device discovery might result in long delays when the operating system boots.
- In certain error conditions with `no_path_retry` or `queue_if_no_path` feature enabled, applications might hang forever. To overcome these conditions, you must enter the following command to all the affected multipath devices: `dmsetup message device 0 "fail_if_no_path"` where device is the multipath device name.
- An I/O hang might occur when a logical drive is unmapped without first deleting the DM device. This limitation applies to only the SLES11 base. It is fixed in later versions.
- Stale entries might be noticed in `multipath -ll` output if the logical drives are unmapped or deleted without first deleting the DM device and its underlying paths. This limitation applies to only the SUSE 11 base OS.
- In device mapper mode select command is issued synchronously for each LUN. With large LUN configurations, slower failovers for DM multipath devices might occur if there is any delay in completing of the mode select command. This limitation applies to only the SUSE 11 base OS. This has been resolved in later versions.
- If the `scsi_dh_rdac` module is not included in `initrd`, slower device discovery might occur, and the `syslog` might get populated with buffer I/O error messages.
- If the storage vendor and model are not included in `scsi_dh_rdac` device handler, slower device discovery might be seen, and the `syslog` might get populated with buffer I/O error messages.
- Use of the DMMP and RDAC failover solutions together on the same host is not supported. Use only one solution at a time.

Linux RDAC (MPP) driver

This section describes how to install the RDAC (MPP) driver for a Linux configuration. ALUA functionality is supported on SLES SP1, RHEL 6.1 or their later versions. IBM recommends using Linux Device Mapper Multipath driver for new servers and storage subsystem configurations. Linux RDAC (MPP) multipath driver support is being discontinued with controller firmware version 7.8x.xx.xx.

Important: Before you install MPP, make sure that the partitions and LUNs are configured and assigned and that the correct HBA driver is installed.

To install MPP, complete the following steps:

1. Download the MPP driver package from the IBM System Storage Disk support portal.
2. Create a directory on the host and download the MPP driver package to that directory.
3. Type the following command to uncompress the file:

```
# tar -zxvf rdac-LINUX-package_version-source.tar.gz
```

where *package_version* is the SLES or RHEL package version number. As a result, a directory called `linuxrdac-version#` or `linuxrdac` is created.
4. Open the readme file that is included in the `linuxrdac-version#` directory.
5. In the readme file, find the instructions for building and installing the driver and complete all of the steps.

Note: Be sure to restart the server before you proceed to the next step.

6. Type the following command to list the installed modules:

```
# lsmod
```

7. Verify that module entries are included in the following `lsmod` list.

Module entries for SLES or RHEL:

- `scsi_mod`
- `sd_mod`
- `sg`
- `mppVhba`
- `mppUpper`
- `lpfc` (or `qla2xxx` for BladeCenter configurations)
- `lpfcdfc` (if `ioctl` module is installed)

Note: If you do not see the `mpp_Vhba` module, the likely cause is that the server was rebooted before the LUNs were assigned, so the `mpp_Vhba` module was not installed. If this is the case, assign the LUNs now, restart the server, and repeat this step.

8. Type the following command to verify the driver version:

```
# mppUtil -V
```

The Linux multipath driver version is displayed.

9. Type the following command to verify that devices are configured with the RDAC driver:

```
# ls -lR /proc/mpp
```

An output similar to the following example is displayed:

```
# ls -lR /proc/mpp
/proc/mpp:
total 0
dr-xr-xr-x  4 root    root          0 Oct 24 02:56 DS4100-sys1
crwxrwxrwx  1 root    root        254,  0 Oct 24 02:56 mppVBusNode

/proc/mpp/ DS4100-sys1:
total 0
dr-xr-xr-x  3 root    root          0 Oct 24 02:56 controllerA
dr-xr-xr-x  3 root    root          0 Oct 24 02:56 controllerB
-rw-r--r--  1 root    root          0 Oct 24 02:56 virtualLun0
-rw-r--r--  1 root    root          0 Oct 24 02:56 virtualLun1
-rw-r--r--  1 root    root          0 Oct 24 02:56 virtualLun2
-rw-r--r--  1 root    root          0 Oct 24 02:56 virtualLun3
-rw-r--r--  1 root    root          0 Oct 24 02:56 virtualLun4
-rw-r--r--  1 root    root          0 Oct 24 02:56 virtualLun5

/proc/mpp/ DS4100-sys1/controllerA:
total 0
dr-xr-xr-x  2 root    root          0 Oct 24 02:56 lpfc_h6c0t2

/proc/mpp/ DS4100-sys1/controllerA/lpfc_h6c0t2:
total 0
-rw-r--r--  1 root    root          0 Oct 24 02:56 LUN0
-rw-r--r--  1 root    root          0 Oct 24 02:56 LUN1
-rw-r--r--  1 root    root          0 Oct 24 02:56 LUN2
-rw-r--r--  1 root    root          0 Oct 24 02:56 LUN3
-rw-r--r--  1 root    root          0 Oct 24 02:56 LUN4
-rw-r--r--  1 root    root          0 Oct 24 02:56 LUN5

/proc/mpp/ DS4100-sys1/controllerB:
total 0
dr-xr-xr-x  2 root    root          0 Oct 24 02:56 lpfc_h5c0t0
```

```

/proc/mpp/ DS4100-sys1/controllerB/lpfc_h5c0t0:
total 0
-rw-r--r-- 1 root root 0 Oct 24 02:56 LUN0
-rw-r--r-- 1 root root 0 Oct 24 02:56 LUN1
-rw-r--r-- 1 root root 0 Oct 24 02:56 LUN2
-rw-r--r-- 1 root root 0 Oct 24 02:56 LUN3
-rw-r--r-- 1 root root 0 Oct 24 02:56 LUN4
-rw-r--r-- 1 root root 0 Oct 24 02:56 LUN5

```

10. Another way to make sure that the RDAC driver discover the available physical logical drives and created virtual logical drives for them, type the command `/opt/mpp/lsvdev`, and press **Enter**. You can now send I/O to the logical drives.
11. If you make any changes to the RDAC configuration file (`/etc/mpp.conf`) or the persistent binding file (`/var/mpp/devicemapping`), run the `mppUpdate` command to rebuild the RAMdisk image to include the new file. In this way, the new configuration file (or persistent binding file) can be used on the next system restart.
12. To dynamically reload the driver stack (`mppUpper`, physical HBA driver modules, `mppVhba`) without restarting the system, perform these steps:
 - a. To unload the `mppVhba` driver, type the command `rmmod mppVhba`, and press **Enter**.
 - b. To unload the physical HBA driver, type the command `modprobe -r "physical hba driver modules"`, and press **Enter**.
 - c. To unload the `mppUpper` driver, type the command `rmmod mppUpper`, and press **Enter**.
 - d. To reload the `mppUpper` driver, type the command `modprobe mppUpper`, and press **Enter**.
 - e. To reload the physical HBA driver, type the command `modprobe "physical hba driver modules"`, and press **Enter**.
 - f. To reload the `mppVhba` driver, type the command `modprobe mppVhba`, and press **Enter**.
13. Restart the system to unload the driver stack.

Note: After you install the RDAC driver, the following commands and pages are available:

- `mppUtil`
- `mppBusRescan`
- `mppUpdate`
- RDAC

About mppUtil

The `mppUtil` utility is a general-purpose command-line driven utility that works only with MPP-based RDAC solutions. The utility instructs RDAC to perform various maintenance tasks but also serves as a troubleshooting tool when necessary. To use the `mppUtil` utility, type the following command, and press **Enter**:

```

mppUtil [-a target_name] [-c wwn_file_name] [-d
debug_level]
[-e error_level] [-g virtual_target_id] [-I host_num]
[-o feature_action_name[=value][, SaveSettings]]
[-s "failback" | "avt" | "busscan" | "forcerebalance"]
[-S] [-U]
[-V] [-w target_wwn,controller_index]

```

Note: The parameters must be in double quotation marks.
The mppUtil utility is a cross-platform tool. The description for each parameter follows.

Table 30. Description of mppUtil parameters

Parameter	Description
<code>-a target_name</code>	Shows the RDAC driver's internal information for the specified virtual <code>target_name</code> (storage subsystem name). If a <code>target_name</code> value is not included, the <code>-a</code> parameter shows information about all of the storage subsystems that are currently detected by this host.
<code>-c wwn_file_name</code>	Clears the WWN file entries. This file is located at <code>/var/mpp</code> with the extension <code>.wwn</code> .
<code>-d debug_level</code>	<p>Sets the current debug reporting level. This option works only if the RDAC driver has been compiled with debugging enabled. Debug reporting is comprised of two segments. The first segment refers to a specific area of functionality, and the second segment refers to the level of reporting within that area. The <code>debug_level</code> is one of these hexadecimal numbers:</p> <ul style="list-style-type: none"> • <code>0x20000000</code>– Shows messages from the RDAC driver's <code>init()</code> routine. • <code>0x10000000</code>– Shows messages from the RDAC driver's <code>attach()</code> routine. • <code>0x08000000</code>– Shows messages from the RDAC driver's <code>ioctl()</code> routine. • <code>0x04000000</code>– Shows messages from the RDAC driver's <code>open()</code> routine. • <code>0x02000000</code>– Shows messages from the RDAC driver's <code>read()</code> routine. • <code>0x01000000</code>– Shows messages related to HBA commands. • <code>0x00800000</code>– Shows messages related to aborted commands. • <code>0x00400000</code>– Shows messages related to panic dumps. • <code>0x00200000</code>– Shows messages related to synchronous I/O activity. • <code>0x00000001</code>– Debug level 1. • <code>0x00000002</code>– Debug level 2. • <code>0x00000004</code>– Debug level 3. • <code>0x00000008</code>– Debug level 4. <p>These options can be combined with the logical AND operator to provide multiple areas and levels of reporting as needed. For use by Customer and Technical Support representatives only.</p>

Table 30. Description of mppUtil parameters (continued)

Parameter	Description
-e <i>error_level</i>	<p>Sets the current error reporting level to <i>error_level</i>, which can have one of these values:</p> <ul style="list-style-type: none"> • 0- Show all errors. • 1- Show path failover, controller failover, repeatable, fatal, and recovered errors. • 2- Show path failover, controller failover, repeatable, and fatal errors. • 3- Show path failover, controller failover, and fatal errors. This is the default setting. • 4- Show controller failover and fatal errors. • 5- Show fatal errors. <p>For use by Customer and Technical Support representatives only.</p>
-g <i>target_id</i>	<p>Displays detailed information about the state of each controller, path, and LUNs for the specified storage subsystem. You can find the <i>target_id</i> by running the <code>dsMUtil -a</code> command.</p>
-M	<p>Shows the MPIO disk-to-drive mappings for the DSM. The output is similar to that found with the SMdevices utility. For use by Customer and Technical Support representatives only.</p>
-o [[<i>feature_action_name</i> [= <i>value</i>]] [<i>feature_variable_name</i> = <i>value</i>]][, SaveSettings]	<p>Troubleshoots a feature or changes a configuration setting. Without the SaveSettings keyword, the changes only affect the in-memory state of the variable. The SaveSettings keyword changes both the in-memory state and the persistent state. Some example commands are:</p> <ul style="list-style-type: none"> • <code>dsMUtil -o</code> Displays all the available feature action names. • <code>dsMUtil -o DisableLunRebalance=0x3 -</code> Turns off the DSM-initiated storage subsystem LUN rebalance (affects only the in-memory state).
-P [GetMpioParameters MpioParameter= <i>value</i> ...]	<p>Displays and sets MPIO parameters. For use by Customer and Technical Support representatives only.</p>
-R	<p>Remove the load-balancing policy settings for inactive devices.</p>

Table 30. Description of mppUtil parameters (continued)

Parameter	Description
-s ["failback" "avt" "busscan" "forcerebalance"]	Manually initiates one of the DSM driver's scan tasks. A "failback" scan causes the DSM driver to reattempt communications with any failed controllers. An "avt" scan causes the DSM driver to check whether AVT has been enabled or disabled for an entire storage subsystem. A "busscan" scan causes the DSM driver to go through its unconfigured devices list to see if any of them have become configured. A "forcerebalance" scan causes the DSM driver to move storage subsystem logical drives to their preferred controller and ignores the value of the DisableLunRebalance configuration parameter of the DSM driver.
-w <i>target_wwn</i> , <i>controller_index</i>	For use by Customer and Technical Support representatives only.

MAC OS multipath driver

The only supported MAC OS multipath driver is the one provided by the HBA vendor - ATTO Corporation. For instructions to install the MAC OS multipath driver, go to the ATTO IBM Solutions website. On that website, see information about installing the Celerity HBA driver and installing the ATTO Configuration Tool in the *ATTO Celerity MultiPaxath Director Installation and Operation Manual*.

Important: After you configure a LUN, you must reboot your server for the LUN to be detected.

Veritas DMP driver

See the Symantec Storage Foundation for Windows documentation for instructions for installing the Veritas DMP driver at <http://www.symantec.com/business/support/>.

Note: The Array Support Library (ASL) that supports DMP on the storage subsystem might have to be loaded. The ASL might be a separate file available from Symantec, or it might be integrated with Volume Manager, depending on the version of Storage Foundation.

HP-UX PV-links

If an HP-UX system is attached with two host bus adapters to the storage subsystem, you can establish redundant access to storage with physical volume links (PV-links), a feature of the HP-UX operating system. PV-links achieve access redundancy with devices that have both primary and secondary paths to the same device.

Important:

- There are two methods for establishing redundant access to storage using PV-links:
 - If you have controller firmware version 07.xx.xx.xx, 06.xx.xx.xx, or 05.xx.xx.xx, use the method described in "Using PV-links: Method 1" on page 144.

- If you have controller firmware version 04.xx.xx.xx, use the method described in “Using PV-links: Method 2” on page 145.
- SMutil must be installed on the host for either method.

Using PV-links: Method 1

For controller firmware version 05.xx.xx.xx or later, use the following procedure to enable multipath I/O with PV-links:

1. Run the **hot_add** command from HP-UX at the shell prompt. This command updates any new devices that are created or added. A dump is generated. When the **hot_add** command runs, each new logical drive that is created in the Subsystem Management window represents a disk device to the operating system.
2. Run the **SMdevices** command. The system provides a dump similar to the example in the table that follows. Notice that every logical drive and logical drive access unit has been assigned a logical unit number (LUN). Each logical drive has two paths. Each storage subsystem controller has one logical drive access.

```
/dev/rds/c166t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0,
Logical Drive WWN <600a0b80000f5d6c00000001e3eaead2b>,
Preferred Path (Controller-B): In Use]
```

```
/dev/rds/c166t0d1 [Storage Subsystem DS4000, Logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>,
Alternate Path (Controller-B): Not In Use]
```

```
/dev/rds/c166t0d2 [Storage Subsystem DS4000, Logical Drive Finance,
LUN 2, Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>,
Alternate Path (Controller-B): Not In Use]
```

```
/dev/rds/c166t0d3 [Storage Subsystem DS4000, Logical Drive Purchasing,
LUN 3, Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>,
Alternate Path (Controller-B): Not In Use]
```

```
/dev/rds/c166t0d4 [Storage Subsystem DS4000, Logical Drive Development,
LUN 4, Logical Drive WWN <600a0b80000f5d6c00000001d3eaeacef>,
Preferred Path (Controller-B): In Use]
```

```
/dev/rds/c166t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f5d6c00000001b0000000>]
```

```
/dev/rds/c172t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0,
Logical Drive WWN <600a0b80000f5d6c00000001e3eaead2b>,
Alternate Path (Controller-A): Not In Use]
```

```
/dev/rds/c172t0d1 [Storage Subsystem DS4000, logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>,
Preferred Path (Controller-A): In Use]
```

```
/dev/rds/c172t0d2 [Storage Subsystem DS4000, Logical Drive Finance, LUN 2,
Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>,
Preferred Path (Controller-A): In Use]
```

```
/dev/rds/c172t0d3 [Storage Subsystem DS4000, Logical Drive Purchasing, LUN 3,
Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>,
Preferred Path (Controller-A): In Use]
```

```
/dev/rds/c172t0d4 [Storage Subsystem DS4000, Logical Drive Development, LUN 4,
Logical Drive WWN <600a0b80000f5d6c00000001d3eaeacef>,
Alternate Path (Controller-A): Not In Use]
```

```
/dev/rds/c172t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f5d6c0000002200000000>]
```

Note: If you do not see the logical drives and logical drive accesses after you run the **hot_add** and **SMdevices** commands, use the **reboot** command to restart the HP-UX host.

3. Determine the preferred and alternate paths for each logical drive by examining the output from the **SMdevices** command, as shown in the previous example. Notice that each device is listed twice; one instance is the preferred path and one instance is the alternate path.

Preferred path

In the following sample output, the preferred path is
/dev/rdisk/c166t0d0.

```
/dev/rdisk/c166t0d0 [Storage Subsystem DS4000, Logical Drive
Accounting, LUN 0, Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>,
Preferred Path (Controller-B): In Use]
```

Alternate path

In the following sample output, the alternate path is
/dev/rdisk/c172t0d0.

```
/dev/rdisk/c172t0d0 [Storage Subsystem DS4000, Logical Drive
Accounting, LUN 0, Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>,
Alternate Path (Controller-A): NotIn Use]
```

Using PV-links: Method 2

For controller firmware version 4.xx.xx.xx, use the following procedures to enable multipath I/O with PV-links:

- Determine the preferred and alternate paths
- Create the logical drives and arrays

Determining preferred and alternate paths

1. Run the **hot_add** command from HP-UX at the shell prompt. This command updates any new devices that are created or added. A dump is generated. When the **hot_add** command runs, each new logical drive that is created in the Subsystem Management window represents a disk device to the operating system.
2. Run the **SMdevices** command. The system provides a dump similar to the following example. Notice that every logical drive and logical drive access unit has been assigned a logical unit number (LUN). Each logical drive has two paths. Each storage subsystem controller has one logical drive access.

```
/dev/rdisk/c166t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0,
Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>]
/dev/rdisk/c166t0d1 [Storage Subsystem DS4000, Logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>]
/dev/rdisk/c166t0d2 [Storage Subsystem DS4000, Logical Drive Finance, LUN 2,
Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>]
/dev/rdisk/c166t0d3 [Storage Subsystem DS4000, Logical Drive Purchasing, LUN 3,
Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>]
/dev/rdisk/c166t0d4 [Storage Subsystem DS4000, Logical Drive Development, LUN 4,
Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>]
/dev/rdisk/c166t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f56d00000001b00000000>]

/dev/rdisk/c172t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0,
Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>]
/dev/rdisk/c172t0d1 [Storage Subsystem DS4000, logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>]
/dev/rdisk/c172t0d2 [Storage Subsystem DS4000, Logical Drive Finance, LUN 2,
Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>]
/dev/rdisk/c172t0d3 [Storage Subsystem DS4000, Logical Drive Purchasing, LUN 3,
Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>]
/dev/rdisk/c172t0d4 [Storage Subsystem DS4000, Logical Drive Development, LUN 4,
Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>]
/dev/rdisk/c172t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f5d6c0000002200000000>]
```

Note: If you do not see the logical drives and logical drive accesses after you run the **hot_add** and **SMdevices** commands, use the **reboot** command to restart the HP-UX host.

3. Determine the preferred and alternate paths for each logical drive by examining the output from the **SMdevices** command, as shown in the example in preceding example.

Notice that each device is listed twice; one instance is the preferred path and one instance is the alternate path. Also, notice that each device has a worldwide name (WWN). Part of the WWN of each logical drive is unique for each controller in the storage subsystem. The WWNs for the logical drive access in the preceding example differ in only five digits, **f56d0** and **f5d6c**.

The devices in the preceding example are viewed through the controllers **c166** and **c172**. To determine the preferred path of a specific logical drive that is seen by the operating system, complete the following steps:

- a. Find the WWN for each logical drive access. In this case, Logical Drive Access 1 is associated with **c166** and has the WWN of **f56d0**.

```
/dev/rdisk/c166t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f56d00000001b00000000>]
```

Logical Drive Access 2 is associated with **c172** and has the WWN of **f5d6c**.

```
/dev/rdisk/c172t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f5d6c0000002200000000>]
```

- b. Identify the preferred device path name for the attached storage device by matching the logical drive WWN to a logical drive access WWN. In this case, the WWN for **LUN 0** is associated with controller **c166** and **c172**. Therefore, the preferred path for **LUN 0** is **/dev/rdsk/c166t0d0**, which is controller **c166**.

```
/dev/rdsk/c166t0d0 [Storage Subsystem DS4000, Logical Drive
Accounting, LUN 0, Logical Drive g<600a0b80000f56d00000001e3eae2b>]
```

The alternate path is **/dev/rdsk/c172t0d0**, which is controller **c172**.

```
/dev/rdsk/c172t0d0 [Storage Subsystem DS4000, Logical Drive
Accounting, LUN 0, Logical Drive WWN <600a0b80000f56d00000001e3eae2b>]
```

- c. To keep a record for future reference, enter this path information for **LUN 0** into a matrix (similar to the one in Table 31).

Table 31. Sample record of logical drive preferred and alternate paths

LUN	Logical drive name	Preferred path	Alternate path
0	Accounting	/dev/rdsk/c166t0d0	/dev/rdsk/c172t0d0
1	HR	/dev/rdsk/c172t0d1	/dev/rdsk/c166t0d1
2	Finance	/dev/rdsk/c172t0d2	/dev/rdsk/c166t0d2
3	Purchasing	/dev/rdsk/c172t0d3	/dev/rdsk/c166t0d3
4	Development	/dev/rdsk/c166t0d4	/dev/rdsk/c172t0d4

- d. Repeat step 3.a through step 3.c for each logical drive that is seen by the operating system.

Creating logical drives and arrays

After you have determined the preferred and alternate paths, and have recorded them in a matrix for future reference, complete the following steps to create logical drives and arrays.

Important: Do not use the HP System Administration Management (SAM) application for DS3000, DS4000, or DS5000 storage configuration; if you do, you might get unwanted results.

Note: The steps in this procedure refer to LUN 0 in Table 31 on page 146.

1. Create a physical logical drive and define the primary paths for the attached storage devices. The primary path is the preferred path. Type the following command at the shell prompt:

```
#pvcreate /dev/rdisk/c166t0d0
```

The system confirms the creation of the new physical logical drive.

2. Create arrays.

Note: For more information about how to create arrays, see the HP-UX documentation or manpages.

- a. Make a directory for the array by typing the following commands. This directory must be in the /dev directory.

```
#cd /dev  
#mkdir /vg1
```

- b. Create the group special file in the /dev directory for the array by typing the following command:

```
#mknod /dev/vg1/group c 64 0x010000
```

- c. Create an array and define physical logical drive names (primary link) for the attached storage device by typing the following command:

```
#vgcreate /dev/vg1/ /dev/dsk/c166t0d0
```

- d. Define the secondary path name (alternate path) for the attached-storage device by typing the following command:

```
#vgextend vg1 /dev/dsk/c172t0d0
```

Note: You can also use the **vgextend** command to add storage devices to an existing array. Add the primary path first, and then add the alternate path, as shown in the following example.

- 1) Add the primary path for LUN1.

```
#vgextend vg1 /dev/dsk/c172t0d1
```

- 2) Add the secondary path for LUN1.

```
#vgextend vg1 /dev/dsk/c166t0d1
```

3. Create logical drives. For more information, see the HP-UX documentation.
4. Create file systems for the logical drives.
5. Repeat step 1 through step 4 to create additional arrays. For more information, see the HP-UX documentation.
6. Verify the primary (preferred) and secondary (alternate) paths for each device by typing the following command:

```
#vgdisplay -v vgname
```

where *vgname* is the array name.

HP-UX native multipathing

Multipathing is native to HP-UX 11i v3. It is built into the mass storage subsystem, and it is available to applications without any special configuration.

For more information about native multipathing in HP-UX 11iv3, see the documentation at <http://docs.hp.com/en/oshpux11iv3.html>.

Solaris failover drivers

A failover driver monitors I/O paths. If a component failure occurs in one of the Fibre Channel paths, the failover driver reroutes all I/O to another path.

Solaris host systems require one of the following failover drivers:

- Solaris Multiplexed I/O (MPxIO)
- RDAC
- Veritas VolumeManager with Dynamic Multipathing (DMP)

Note:

1. RDAC is not supported on Solaris 10. You must use either Solaris MPxIO or the Veritas DMP failover driver.
2. With Solaris 10, MPxIO capability is built in. If you want to use MPxIO with previous versions of Solaris, you must install SUN StorEdge SAN Foundation Suite.
3. To support Asymmetric Logical Unit Access (ALUA), you must have Solaris 11 or later.

This section includes the following procedures:

- “Installing the MPxIO driver”
- “Installing the RDAC failover driver on Solaris and modifying the configuration files” on page 154
- “Planning and preparing for Veritas DMP driver installation” on page 156

Installing the MPxIO driver

Multiplexed I/O (MPxIO) is a Sun Solaris multipath driver architecture. This failover driver enables storage subsystems to be accessed through multiple host controller interfaces from a single instance of the storage subsystem. MPxIO helps protect against storage subsystem outages that are caused by controller failures. If one controller fails, MPxIO automatically switches to an alternate controller.

MPxIO is fully integrated within the Solaris 10 operating system. For Solaris 8 and 9 operating systems, MPxIO is available as part of the Sun StorEdge SAN Foundation Suite and must be installed separately.

For the latest supported version of Sun StorEdge SAN Foundation Suite, the latest Solaris kernel patches, and the most recent updates to information about using MPxIO, see the Storage Manager readme file for Solaris. See “Finding Storage Manager software, controller firmware, and readme files” on page xiv for steps to find the readme file on the web.

This section contains the following topics:

- “Device name change considerations for MPxIO” on page 149
- “Acquiring the latest MPxIO driver version” on page 149
- “Enabling the MPxIO failover driver” on page 149
- “Disabling the MPxIO multipath driver” on page 154

Note: For more information, see the *Sun StorEdge SAN Foundation Software Installation Guide*, *Sun StorEdge SAN Foundation Software Configuration Guide*, and *Sun Solaris Fibre Channel and Storage Multipathing Administration Guide* on the Sun website <http://docs.sun.com>.

Device name change considerations for MPxIO:

In the `/dev` and `/devices` trees, devices are named differently from their original names when MPxIO is enabled. See the following examples.

Device name with MPxIO disabled:

```
/dev/dsk/c1t1d0s0
```

Device name with MPxIO enabled:

```
/dev/rdisk/c0t600A0B800011121800006B31452CC6A0d0s2
```

You must configure the applications that engage the device directly to use the new names whenever the MPxIO configuration is enabled or disabled.

In addition, the `/etc/vfstab` file and the dump configuration also contain references to device names. When you use the **stmsboot** command to enable or disable MPxIO, as described in the following sections, `/etc/vfstab` and the dump configuration are automatically updated with the new device names.

Acquiring the latest MPxIO driver version: The method of acquiring MPxIO depends upon which version of Solaris you have installed:

Solaris 10

MPxIO is integrated within the Solaris 10 operating system and does not have to be installed separately. Use Solaris 10 patches to update MPxIO with Solaris 10. The patches are available at the Sun Technical Support website <http://sunsolve.sun.com>.

Note: You must install the regular kernel jumbo patch, because there are dependencies between the various patches that make up the driver stack.

Solaris 8 and 9

Because MPxIO is not included with Solaris 8 and 9, you must download the required SAN suite (Sun StorEdge SAN Foundation Suite) from the Sun Technical Support website <http://sunsolve.sun.com>. On this webpage, click **SAN 4.4 release Software/Firmware Upgrades & Documentation**.

Note: Use the `install_it.ksh` script that is provided to install the software.

Enabling the MPxIO failover driver: This section describes how to enable MPxIO by using the **stmsboot** command. In addition to enabling MPxIO, this command also updates the device names in the `/etc/vfstab` file and the dump configuration files during the next reboot.

Note: In Solaris 10, the **stmsboot** command is used to enable or disable MPxIO on all devices.

Before you begin:

1. Install the Solaris operating system, and the latest patches.
2. Make sure that the Solaris host type was selected when the host was defined.

Enabling MPxIO on Solaris 8 and 9

1. Install the latest version of Sun StorEdge SAN Foundation Suite and required patches, using the **Sun StorEdge install_it** script. For more information, see the *Sun StorEdge SAN Foundation Suite x.xx Installation Guide* (where *x.xx* is the version of the StorEdge software).

2. Edit the `/kernel/drv/scsi_vhci.conf` configuration file to make sure that the VID/PID is not specified in this file. Also, make sure that the following entries are in the file:

```
mpxio-disable="no";  
load-balance="none";  
auto-failback="enable";
```

Note: In a cluster environment where logical drives (LUNs) are shared between multiple Sun servers, you might have to set the `auto-failback` parameter to `disable` to prevent the following phenomenon, which can occur when one of the servers has a failed path to one of the shared LUNs.

If a host in a cluster server configuration loses a physical path to a storage subsystem controller, LUNs that are mapped to the cluster group can periodically failover and then failback between cluster nodes until the failed path is restored. This behavior is the result of the automatic logical drive failback feature of the multipath driver. The cluster node with a failed path to a storage subsystem controller issues a failover command for all LUNs that were mapped to the cluster group to the controller that it can access. After a programmed interval, the nodes that did not have a failed path issue a failback command for the LUNs because they can access the LUNs on both controllers. The cluster node with the failed path is unable to access certain LUNs. This cluster node then issues a failover command for all LUNs, which repeats the LUN failover-failback cycle.

For supported cluster services, see the System Storage Interoperation Center at www.ibm.com/systems/support/storage/config/ssic

3. If you made any changes to the `/kernel/drv/scsi_vhci.conf` file in the previous step, save the file and use the following command to restart the server:

```
# shutdown -g0 -y -i6
```
4. If necessary, update the Fibre Channel HBA firmware.
5. Create the storage subsystem logical drives and map them to the Fibre Channel HBA ports in the Sun servers.

Enabling MPxIO on Solaris 10

Note the following considerations for `stmsboot -e [enable] -d [disable] and -u [update]`:

- When you run the **stmsboot** command, accept the default to Reboot the system now.
- The **stmsboot** command saves copies of the original `/kernel/drv/fp.conf` and `/etc/vfstab` files before it modifies them, so you can use the saved files to recover from any unexpected problems.
- Make sure that the **eeeprom** boot device is set to boot from the current boot device.

To enable MPxIO on all Fibre Channel devices, complete the following steps:

1. Run the **stmsboot -e** command, and select the default `[y]` to restart the server:

```
# stmsboot -e
```

```
WARNING: This operation will require a reboot.  
Do you want to continue ? [y/n] (default: y) y  
The changes will come into effect after rebooting the system.  
Reboot the system now ? [y/n] (default: y) y
```


Note: During the reboot, /etc/vfstab and the dump configuration are updated to reflect the device name changes.

2. After the reboot, configure your applications to use new device names, as explained in “Device name change considerations for MPxIO” on page 149.
3. If necessary, edit the /kernel/drv/fp.conf configuration file to verify that the following parameter is set as follows:

```
mpxio-disable="no";
```

Edit the /kernel/drv/scsi_vhci.conf configuration file to verify that the following parameters are set as follows:

```
load-balance="none";
auto-failback="enable";
```

4. If you made any changes to configuration files in the previous step, save the file, and use the following command to restart the server:

```
# shutdown -g0 -y -i6
```
5. If necessary, update the Fibre Channel HBA firmware.
6. Create the storage subsystem logical drives and map them to the Fibre Channel HBA ports in the Sun servers.

Verifying devices and configuring failover/failback path for the mapped LUNs

To verify devices and configure the failover path for the mapped LUNs, complete the following steps:

1. Use the **cfgadm -al** command to verify the devices. Run the **cfgadm -al** command to display information about the host ports and their attached devices:

```
# cfgadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
PCI0	vgs8514/hp	connected	configured	ok
PCI1	unknown	empty	unconfigured	unknown
PCI2	unknown	empty	unconfigured	unknown
PCI3	mult/hp	connected	configured	ok
PCI4	unknown	empty	unconfigured	unknown
PCI5	unknown	empty	unconfigured	unknown
PCI6	unknown	empty	unconfigured	unknown
PCI7	mult/hp	connected	configured	ok
PCI8	mult/hp	connected	configured	ok
c0	scsi-bus	connected	configured	unknown
c0::dsk/c0t6d0	CD-ROM	connected	configured	unknown
c1	fc-private	connected	configured	unknown
c1::500000e0106fca91	disk	connected	configured	unknown
c1::500000e0106fcde1	disk	connected	configured	unknown
c1::500000e0106fcf31	disk	connected	configured	unknown
c1::500000e0106fd061	disk	connected	configured	unknown
c1::500000e0106fd7b1	disk	connected	configured	unknown
c1::500000e0106fdaa1	disk	connected	configured	unknown
c1::50800200001d9841	ESI	connected	configured	unknown
c2	fc-fabric	connected	configured	unknown
c2::201400a0b811804a	disk	connected	configured	unusable
c2::201400a0b8118098	disk	connected	configured	unusable
c2::201700a0b8111580	disk	connected	configured	unusable
c3	fc-fabric	connected	configured	unknown
c3::201500a0b8118098	disk	connected	configured	unusable
c3::201600a0b8111580	disk	connected	configured	unusable
c3::202500a0b811804a	disk	connected	configured	unusable
c4	fc-fabric	connected	configured	unknown
c4::200400a0b80f1285	disk	connected	configured	unknown
c4::200400a0b8127a26	disk	connected	configured	unusable
c5	fc-fabric	connected	configured	unknown

c5::200400a0b82643f5	disk	connected	unconfigured	unknown
c5::200500a0b80f1285	disk	connected	configured	unknown
c5::200500a0b8127a26	disk	connected	configured	unusable
c5::200c00a0b812dc5a	disk	connected	configured	unknown
usb0/1	usb-kbd	connected	configured	ok
usb0/2	usb-mouse	connected	configured	ok
usb0/3	unknown	empty	unconfigured	ok
usb0/4	unknown	empty	unconfigured	ok
#				

- You can also display information about the attachment points on a server. In the following example, c0 represents a fabric-connected host port, and c1 represents a private, loop-connected host port. Use the **cfgadm** command to manage the device configuration on fabric-connected host ports. By default, the device configuration on private, loop-connected host ports is managed by the Solaris host.

Note: The **cfgadm -l** command displays information about Fibre Channel host ports. Also use the **cfgadm -al** command to display information about Fibre Channel devices. The lines that include a port World Wide Name (WWN) in the Ap_Id field associated with c0 represent a fabric device. Use the **cfgadm configure** and **cfgadm unconfigure** commands to manage those devices and make them available to Solaris hosts.

```
# cfgadm -l
Ap_Id      Type      Receptacle  Occupant    Condition
c0         fc-fabric connected   unconfigured unknown
c1         fc-private connected   configured  unknown
```

- Configure the device using the following command:

```
cfgadm -c configure Ap-Id
```

The **Ap_ID** parameter specifies the attachment point ID of the configured Fibre Channel devices. This ID can be the controller number and WWN of a device (for example, c3::50020f230000591d).

See the output example in step 1. Also, see the *cfgadm* manpage for an explanation of attachment points.

Note: An Ap_Id with type fc-private cannot be unconfigured. Only the type fc-fabric can be configured and unconfigured.

- Use the **luxadm probe** command to list all mapped LUNs:

```
# luxadm probe
luxadm probe
No Network Array enclosures found in /dev/es

Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006ADE452CBC62d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006ADF452CBC6Ed0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AE0452CBC7Ad0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AE1452CBC88d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AE2452CBC94d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AE3452CBCA0d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AE4452CBCACd0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AE5452CBCB8d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AE6452CBC4d0s2
```

```

Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AE7452CBCD2d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AE8452CBCEd0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AE9452CBCEAd0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AEA452CBCF8d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AEB452CBD04d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AEC452CBD10d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006AED452CBD1Ed0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006B2A452CC65Cd0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006B2B452CC666d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006B2C452CC670d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006B2D452CC67Ad0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006B31452CC6A0d0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c0t600A0B800011121800006B32452CC6ACd0s2
Node WWN:200400a0b8111218 Device Type:Disk device
Logical Path:/dev/rdisk/c8t201400A0B8111218d7s2

```

5. You can use the **luxadm display logical path** command to list more details about each mapped LUN, including the number of paths to each LUN. The following example uses a logical path from the previous example.

```

# luxadm display /dev/rdisk/c0t600A0B800011121800006B31452CC6A0d0s2
DEVICE PROPERTIES for disk: /dev/rdisk/c0t600A0B800011121800006B31452CC6A0d0s2
Vendor: IBM
Product ID: 1742-900
Revision: 0914
Serial Num: 1T51207691
Uninitialized capacity: 1024.000 MBytes
Write Cache: Enabled
Read Cache: Enabled
Minimum prefetch: 0x0
Maximum prefetch: 0x0
Device Type: Disk device
Path(s):

/dev/rdisk/c0t600A0B800011121800006B31452CC6A0d0s2
/devices/scsi_vhci/ssd@g600a0b800011121800006b31452cc6a0:c,raw
Controller /devices/pci@7c0/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0
Device Address 201400a0b8111218,1e
Host controller port WWN 210100e08ba0fca0
Class secondary
State STANDBY
Controller /devices/pci@7c0/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0
Device Address 201500a0b8111218,1e
Host controller port WWN 210100e08ba0fca0
Class primary
State ONLINE
Controller /devices/pci@7c0/pci@0/pci@8/SUNW,qlc@0/fp@0,0
Device Address 201400a0b8111218,1e
Host controller port WWN 210000e08b80fca0
Class secondary
State STANDBY
Controller /devices/pci@7c0/pci@0/pci@8/SUNW,qlc@0/fp@0,0
Device Address 201500a0b8111218,1e

```

```

Host controller port WWN    210000e08b80fca0
Class                      primary
State                      ONLINE
#

```

Unconfiguring a failover/failback path

To unconfigure a failover/failback path, complete the following steps:

1. Before you unconfigure a fabric device, stop all activity to the device and unmount any file systems on the fabric device (see the Solaris administration documentation for unmounting procedures).
2. Run the **cfgadm -al** command to display information about the host ports and their attached devices.
3. Unconfigure the LUN by running the following command:
`cfgadm -c unconfigure Ap-Id`

Where Ap-Id is the LUN that you want to unconfigure.

4. Run the **cfgadm -al** command again to verify that the LUN is now unconfigured.
5. If necessary, define the file structure using the **newfs** command. Add entries to the `/etc/vfstab` file.
6. Type the following command to restart the server:
`shutdown -g0 -y -i6`

Disabling the MPxIO multipath driver: To disable the MPxIO multipath driver, take the applicable action for your version of Solaris:

- For Solaris 10, unconfigure all devices by using the `cfgadm -c unconfigure AP-id Ap-id` command. Then, run the `stmsboot -d` command, and accept the default to Reboot the system now.
- For Solaris 8 and 9, unconfigure all devices by using the `cfgadm -c unconfigure AP-id Ap-id` command, and edit the `/kernel/drv/scsi_vhci.conf` configuration file to set the value of the `mpxio-disable` parameter to `yes`. Restart the server.

To learn how to revert the patches or use the StorEdge software, see the *Sun StorEdge SAN Foundation Installation Software Guide* at <http://docs.sun.com>.

Installing the RDAC failover driver on Solaris and modifying the configuration files

This section describes how to install RDAC on a Solaris host.

Before you begin:

1. RDAC is supported only on Solaris 8 and 9.
2. Because you cannot run both RDAC and MPxIO, make sure that MPxIO is disabled. Check the configuration files (`/kernel/drv/scsi_vhci.conf`, `/kernel/drv/fp.conf`, or both) and make sure that the value of the `mpxio-disable` parameter is set to `Yes`.
3. You must install an HBA driver package before you install RDAC. If you have a SAN-attached configuration, you must also modify the HBA configuration file before you install RDAC. If you fail to follow the procedures in this order, problems can occur.
4. If you modify the failover settings in the HBA configuration file after you install RDAC, you must remove the RDAC from the host.

Important: In some configurations, a patch is required for RDAC to function correctly. Before you begin the RDAC installation, check the Storage Manager readme file for Solaris to find out whether the patch is required for your specific configuration. You can also find the latest RDAC versions and other important information in the readme file. For more information about how to find the readme file on the web, see “Finding Storage Manager software, controller firmware, and readme files” on page xiv.

Installing the RDAC failover driver

To install the RDAC failover driver, complete the following steps:

1. Insert the Solaris installation DVD in the DVD drive.

Note: In this procedure, the installation DVD is mounted at /cdrom/SM91. Modify these commands as needed for your installation. It can be 'cdrom' or 'dvdrom' depending on the way the disk has been mounted.

2. Type the following command to start installing the RDAC package:

```
# pkgadd -d path/filename.pkg
```

where *path/filename* is the directory path and name of the package that you want to install.

The installation process begins.

Information about the packages that can be installed in the specified directory is displayed on the command line, as in the following example:

The following packages are available:

```
1 RDAC                                Redundant Disk Array Controller
                                     (sparc) version number
```

Select package(s) you wish to process (or 'all' to process all packages). (default:all) [?,??,q]:

3. Type the value of the package that you are installing and press Enter. The installation process begins.
4. The software automatically checks for package conflicts. If any conflicts are detected, a message is displayed that indicates that some files are already installed and are in use by another package. The following prompt is displayed:

Do you want to install these conflicting files [y, n, ?]

Type y and press Enter. The following prompt is displayed:

This package contains scripts which will be executed with super-user permission during the process of installing this package.

Do you want to continue with the installation of <RDAC>

[y, n, ?]

5. Type y and press Enter. The installation process continues. When the RDAC package is successfully installed, the following message is displayed:

Installation of <RDAC> was successful.

6. Make sure that the variables in the configuration files for the JNI adapter cards have been set to the correct values.

7. Type the following command to restart the Solaris host:

```
# shutdown -g0 -y -i6
```

Modifying the configuration files

Important: Any modification to the persistent bindings in the `jnic146x.conf` configuration file requires the removal of RDAC. After the RDAC is removed, you can modify the persistent bindings in the `jnic146x.conf` file and then reinstall RDAC.

Complete the following steps to modify the `sd.conf` or `jnic146x.conf` files:

1. Type the following command to remove RDAC:

```
# pkgrm RDAC_driver_pkg_name
```

where `RDAC_driver_pkg_name` is the name of the RDAC driver package that you want to remove.

2. Type the following command to verify that the RDAC drive package is removed:

```
# pkginfo RDAC_driver_pkg_name
```

where `RDAC_driver_pkg_name` is the name of the RDAC driver package that you removed.

3. Type the following command to restart the Solaris host:

```
# shutdown -g0 -y -i6
```

4. Type the following command to modify persistent bindings in the `sd.conf` file or edit the `sd.conf` file:

```
# vi /kernel/drv/jnic146x.conf or sd.conf
```

5. After you have finished making changes, type the following command to save the changes:

```
# :wq
```

6. Type the following command to install the RDAC driver package:

```
# pkgadd -d RDAC_driver_pkg_name
```

where `RDAC_driver_pkg_name` is the name of the RDAC driver package that you want to install.

7. Type the following command to verify package installation:

```
# pkginfo RDAC_driver_pkg_name
```

where `RDAC_driver_pkg_name` is the name of the RDAC driver package that you installed.

8. Type the following command to restart the Solaris host:

```
# shutdown -g0 -y -i6
```

Note: You must restart the host after you modify the `jnic146x.conf` file, because the `jnic146x.conf` driver is read only during the boot process. Failure to restart the host might cause some devices to be inaccessible.

Planning and preparing for Veritas DMP driver installation

This section describes how to plan and prepare for Veritas Dynamic Multipathing (DMP) driver installation. The Veritas DMP driver is a failover driver for Solaris hosts. The DMP failover driver is a feature of Veritas Volume Manager, which is a component of the Storage Foundation product from Symantec. Although RDAC allows you to have only 32 LUNs, DMP allows you to have up to 256 LUNs.

Important: This document does not describe how to install the Veritas product. For more information, see the Symantec documentation at <http://www.symantec.com/business/support/>.

System requirements

Make sure that your server meets the following requirements for installing Veritas DMP:

- Solaris operating system
- Veritas Volume Manager 4.0, 4.1, 5.0, or 5.1
- Array Support Library (ASL), which enables Solaris to recognize the DS3000, DS4000, or DS5000 machine type

Note: The ASL might be a separate file that is available from Symantec, or it might be integrated with Volume Manager, depending on the version of Storage Foundation.

Veritas DMP installation overview

Make sure that your server meets the following prerequisites for installing Veritas DMP:

- The HBAs are installed on the Solaris host.
- The parameter settings in the HBA configuration file (for example, `qla2300.conf`) are modified.
- In a SAN environment, bindings are configured.
- The zones are created and enabled for the Solaris partition.
- Storage is mapped to the Solaris partition.

Complete the procedures in the following order to complete the Veritas DMP installation:

1. "Preparing for Veritas DMP installation"
2. "Installing the ASL package" on page 159

Preparing for Veritas DMP installation

To prepare the host for Veritas DMP installation, complete the following steps:

1. Choose the Solaris host on which you want to install DMP.
2. Complete the following steps to manually define the targets and LUNs in the `/kernel/drv/sd.conf` file:

Note: By default, the `/kernel/drv/sd.conf` file defines targets 0, 1, 2, and 3. LUN0 also is defined for targets 0, 1, 2, and 3. Each target represents a controller to a storage subsystem, and each LUN represents a logical drive. If you are adding target or LUN definitions to the `/kernel/drv/sd.conf` file for an existing DMP configuration, be sure to reboot the Solaris host.

- a. Type the following command to open the `/kernel/drv/sd.conf` file with the vi Editor:

```
# vi /kernel/drv/sd.conf
```

The file looks similar to the following example:

```
#  
# Copyright (c) 1992, Sun Microsystems, Inc.  
#
```

```
# ident "@(#)sd.conf 1.9 98/01/11 SMI"

name="sd" class="scsi" class_prop="atapi"
target=0 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=1 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=2 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=3 lun=0;
```

- b. Use the vi Editor to add target and LUN definitions. In the following example, it is assumed that the Solaris host is attached to one storage subsystem with three LUNs mapped to the subsystem storage partition. In addition, the access LUN must be mapped to the partition.

```
#
# Copyright (c) 1992, Sun Microsystems, Inc.
#
# ident "@(#)sd.conf 1.9 98/01/11 SMI"

name="sd" class="scsi" class_prop="atapi"
target=0 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=1 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=2 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=3 lun=0;

name="sd" class="scsi" target=0 lun=1;
name="sd" class="scsi" target=0 lun=2;
name="sd" class="scsi" target=0 lun=3;
name="sd" class="scsi" target=0 lun=31;
name="sd" class="scsi" target=1 lun=1;
name="sd" class="scsi" target=1 lun=2;
name="sd" class="scsi" target=1 lun=3;
name="sd" class="scsi" target=1 lun=31;
```

- c. Type the following command to save the new entries in the /kernel/drv/sd.conf file:

```
# :wq
```

3. Type the following command to verify that RDAC is not installed on the host:


```
# pkginfo -l RDAC
```
4. If RDAC is installed, type the following command to remove it:


```
# pkgrm RDAC
```
5. Verify that a host partition has been created.

Important: Set the host type to **Solaris with DMP**. Failure to do so results in an inability to map for more than the RDAC limit of 32 LUNs and causes other unwanted results.

6. Make sure that all of the paths are optimal and are in a preferred path state from the SMclient.
7. Install the Veritas Storage Foundation Solaris, including Veritas Volume Manager with DMP.

Important: Before you install Veritas Storage Foundation Solaris with Veritas Volume Manager and DMP, make sure that you have the required license keys. This document does not describe how to install the Veritas product. For more information, see the Symantec documentation at <http://www.symantec.com/business/support/>.

8. Type the following command to restart the Solaris host:

```
# shutdown -g0 -y -i6
```

Installing the ASL package

If required, complete the following steps to install the ASL package.

Note: The VxVM 4.x version of the ASL package is named **SMibmasl** (for more information, see <http://seer.entsupport.symantec.com/docs/284913.htm>). For VxVM version 5.0 and later, many ASLs are integrated into VxVM and do not have to be installed. For VxVM version 5.0 and later, the ASL package is named **VRTSLSIas1** (see <http://seer.entsupport.symantec.com/docs/340469.htm>). The following example assumes that VxVM 4.x is being installed.

1. Type the following command to install the **SMibmasl** package:

Note: You can either select the default (all) or select option 1.

```
# pkgadd -d SMibmasl_pkg
```

2. Type the following command to restart the Solaris host:

```
# shutdown -g0 -y -i6
```

See the Symantec Veritas documentation for information about how to complete the following tasks:

- Start the Veritas Volume Manager
- Set up disk groups
- Create volumes
- Create file systems
- Mount file systems

Identifying devices

After you have installed the multipath driver or verified that the multipath driver is already installed, use the **SMdevices** utility to identify a storage subsystem logical drive that is associated with an operating-system device.

Using the SMdevices utility

The **SMutil** software includes a utility called **SMdevices** that you can use to view the storage subsystem logical drive that is associated with a particular operating-system device name. This utility is helpful when you want to create drive letters or partitions by using Disk Administrator.

Using SMdevices on Windows operating systems

After you create the logical drives on a storage subsystem, go to the host that is attached to that storage subsystem, and complete the following steps to use **SMdevices** on Windows:

1. From a DOS or command prompt, change to the directory
<installation_directory>\Util,

where *installation_directory* is the directory in which you installed the SMutil.

The default directory is c:\Program Files\IBM_DS4000\Util.

2. Type SMdevices and press Enter.

Using SMdevices on UNIX-type operating systems

You can use SMdevices to map the host-assigned device name for each LUN back to its corresponding storage-subsystem device. In the SMdevices output, you can view the following storage-subsystem information, as it is shown on SMclient.

Note: The examples in the list refer to the sample SMdevices output.

- Host assigned name (/dev/sdh)
- DS3000, DS4000, or DS5000 storage subsystem name (DS4500_Storage_Server-A)
- Logical drive name (Raid-5-0A)
- LUN ID (LUN 4)
- Preferred controller owner, and whether that controller is controlling the logical drive

The following example shows a sample SMdevices output for the DS4500_Storage_Server-A storage subsystem:

```
# SMdevices
IBM FASTT Storage Manager Devices, Version 09.12.A5.00
Built Fri Jan 14 16:42:15 CST 2005
(C) Copyright International Business Machines Corporation,
2004 Licensed Material - Program Property of IBM. All rights reserved.

/dev/sdh (/dev/sg10) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-5-0A, LUN 4, Logical Drive ID
<600a0b80000f0fc300000044412e2dbf>, Preferred Path (Controller-A): In Use]
/dev/sdd (/dev/sg6) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-5-1A, LUN 0, Logical Drive ID
<600a0b80000f13ec00000016412e2e86>, Preferred Path (Controller-B): In Use]
/dev/sde (/dev/sg7) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-0-0A, LUN 1, Logical Drive ID
<600a0b80000f0fc30000003c412e2d59>, Preferred Path (Controller-A): In Use]
/dev/sdf (/dev/sg8) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-1-0A, LUN 2, Logical Drive ID
<600a0b80000f0fc30000003e412e2d79>, Preferred Path (Controller-A): In Use]
/dev/sdg (/dev/sg9) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-3-0A, LUN 3, Logical Drive ID
<600a0b80000f13ec00000012412e2e4c>, Preferred Path (Controller-A): In Use]
```

Identifying devices on AIX hosts

The information in this section describes device discovery on AIX. For troubleshooting information about disk array errors on AIX, see “Resolving disk array errors on AIX” on page 266 in Chapter 7, “Troubleshooting,” on page 245.

Understanding devices on AIX hosts

The multipath driver creates the following devices that represent the storage subsystem configuration:

- dar** The disk array router (dar) device represents the entire array, including the current and the deferred paths to all LUNs (hdisks).
- dac** The disk array controller (dac) devices represent a controller within the storage subsystem. There are two dacs in the storage subsystem. With MPIO, the dac device is shown only if a UTM device is assigned.

hdisk Each hdisk device represents an individual LUN on the array.

utm The universal transport mechanism (utm) device is used only with in-band management configurations, as a communication channel between the SMagent and the storage subsystem.

Note: The utm device might be listed in command output, regardless of whether you have an in-band management configuration. For example, a utm might be listed when you run the **lsattr** command on a dac.

Performing the initial device discovery

To perform the initial device discovery, complete the following steps:

1. Make sure that the storage subsystem has been set up, LUNs have been assigned to the host, and the multipath driver has been installed.

2. Type the following command to probe for the new devices:

```
# cfgmgr -v
```

Note: In a SAN configuration, the devices do not log in to the SAN switch until you run the **cfgmgr** command.

3. Type the following command:

```
# lsdev -Cc disk
```

4. Examine the output of the **lsdev -Cc disk** command to make sure that the RDAC software recognizes the storage subsystem logical drives, as shown in the following list:

- Each DS4200 logical drive is recognized as an 1814 DS4200 Disk Array Device.
- Each DS4300 logical drive is recognized as an 1722-600 (600) Disk Array Device.
- Each DS4400 logical drive is recognized as an 1742-700 (700) Disk Array Device.
- Each DS4500 logical drive is recognized as an 1742-900 (900) Disk Array Device.
- Each DS4700 logical drive is recognized as an 1814 DS4700 Disk Array Device.
- Each DS4800 logical drive is recognized as an 1815 DS4800 Disk Array Device.

Important: You might discover that the configuration process has created two dacs and two dars on one storage subsystem. This situation can occur when the host is using a partition that does not have any associated LUNs. When that happens, the system cannot associate the two dacs under the correct dar. If there are no LUNs, the system generates two dacs as expected, but it also generates two dars.

The following list shows the most common causes:

- You create a partition and attach the LUNs to it, but you do not add the host ports to the partition. Therefore, the host ports remain in the default partition.
- You replace one or more HBAs but do not update the worldwide name (WWN) of the partition for the HBA.
- You switch the storage subsystem from one set of HBAs to another as part of a reconfiguration and do not update the WWNs.

In each of these cases, resolve the problem, and run **cfgmgr** again. The system removes the extra dar or moves it from the Available state to the Defined state. (If the system moves the dar into the Defined state, you can delete it.)

Note: When you perform the initial device identification, the Object Data Manager (ODM) attributes of each device are updated with default values. In most cases and for most configurations, the default values are satisfactory. However, there are some values that can be modified for maximum performance and availability. See Appendix D, “Viewing and setting AIX Object Data Manager (ODM) attributes,” on page 301 for information about using the **lsattr** command to view attribute settings on an AIX system.

Example of an initial discovery with MPIO

The following example shows an initial discovery with MPIO.

```
# lsdev -C |grep hdisk10
hdisk10    Available 05-08-02      MPIO Other DS4K Array Disk

# lscfg -vpl hdisk10
hdisk10    U787F.001.DPM0H2M-P1-C3-T1-W200400A0B8112AE4-L9000000000000
MPIO Other DS4K Array Disk
  Manufacturer.....IBM
  Machine Type and Model.....1814      FAStT
  ROS Level and ID.....30393136
  Serial Number.....
  Device Specific.(Z0).....0000053245004032
  Device Specific.(Z1).....

# mpio_get_config -A
Storage Subsystem worldwide name: 60ab8001122ae000045f7fe33
Storage Subsystem Name = 'Kinks-DS-4700'
  hdisk      LUN #
  hdisk2      1
  hdisk3      2
  hdisk4      3
  hdisk5      4
  hdisk6      5
  hdisk7      6
  hdisk8      7
  hdisk9      8
  hdisk10     9
  hdisk11     10
```

Configuring devices

To maximize your storage subsystem performance, you can set the queue depth for your hdisks, disable cache mirroring, use dynamic capacity and dynamic logical drive expansion (DVE), and check the size of your LUNs.

Using the hot_add utility

The **hot_add** utility enables you to add new logical drives without restarting the system. The utility registers the new logical drives with the operating system so that you can use Disk Administrator to create partitions and add device names. The **hot_add** utility is part of the SMutil software package. If you run the program twice and the new logical drives are not displayed in the Disk Administrator window, you must either run Fibre Channel diagnostics or restart the host.

After you create logical drives on a particular storage subsystem, go to the host that is attached to that storage subsystem and complete the following steps to use the **hot_add** utility:

1. From a DOS or command prompt, change to the following directory:
`<installation_directory>\Util`
where *installation_directory* is the directory in which you installed the SMutil.

Note: The default directory is `c:\Program Files\IBM_DS4000\Util`.

2. From a DOS or command prompt, type the following command:
`hot_add`
3. Press Enter. The new logical drives are available through the Disk Administrator.

Using the SMrepassist utility

Use the SMrepassist utility to flush cached data for a logical drive.

Important: The FlashCopy drive cannot be added or mapped to the same server that has the base logical drive of the FlashCopy logical drive in a Windows 2000, Windows Server 2003, Windows Server 2008, or NetWare environment. You must map the FlashCopy logical drive to another server.

To flush cached data in a logical drive, complete the following steps:

1. From a DOS or command prompt, change to the directory
`<installation_directory>\Util`
where *installation_directory* is the directory in which you installed the SMutil.

Note: The default directory is `c:\Program Files\IBM_DS4000\Util`.

2. Type the following command:
`smrepassist -f logical_drive_letter:`
where *logical_drive_letter* is the operating-system drive letter that was assigned to the disk partition on the logical drive.
3. Press Enter.

Stopping and restarting the host-agent software

You must stop and restart the host-agent software if you add storage subsystems to the management domain of the host-agent software. When you restart the service, the host-agent software discovers the new storage subsystems and adds them to the management domain.

Note: If none of the access logical drives are detected after a restart, the host-agent software stops running automatically. Make sure that there is a Fibre Channel connection from the host to the SAN to which the storage subsystem is connected, and restart the host or cluster node so that new host-agent-managed storage subsystems can be discovered.

Use the applicable procedure in this section for your operating system.

Windows 2000

To stop and restart the host-agent software on Windows 2000, complete the following steps:

1. Click **Start > Programs > Administrative Tools > Services**. The Services window opens.

2. Right-click **IBM DS Storage Manager Agent**.
3. Click **Restart**. The Storage Manager Agent stops and then starts again.
4. Close the Services window.

Windows Server 2003 and 2008

To stop and restart the host-agent software on Windows Server 2003 and 2008, complete the following steps:

1. Click **Start > Administrative Tools > Services**. The Services window opens.
2. Right-click **IBM DS Storage Manager Agent**.
3. Click **Restart**. The Storage Manager Agent stops and then starts again.
4. Close the Services window.

Windows Server 2012

To stop and restart the host-agent software on Windows Server 2012, complete the following steps:

1. Click **Server Manager > Tools > Services**. The Services window opens.
2. Right-click **IBM DS Storage Manager Agent**.
3. Click **Restart**. The Storage Manager Agent stops and then starts again.
4. Close the Services window.

Setting the queue depth for hdisk devices

Setting the `queue_depth` attribute to the appropriate value is important for optimal system performance. Use this setting if you have a large storage-subsystem configuration with many attached logical drives and hosts.

This section provides methods for calculating maximum queue depth, which you can use as a guideline to help you determine the best queue-depth setting for your configuration.

Calculating maximum queue depth

The formula for calculating the maximum queue depth for your system depends on which firmware version is installed on the controller. Use one of the following formulas to calculate the maximum queue depth for your system.

Important:

1. The maximum queue depth might not be an optimal setting in all cases. Use the maximum queue depth as a guideline, and adjust the setting as necessary for your specific configuration.
2. In systems with one or more attached SATA devices, you might have to set the `queue_depth` attribute to a lower value than the maximum queue depth.

Formulas for controller firmware version 07.10.xx.xx and later

On DS4800 and DS4700 or DS4200 storage systems that are running storage subsystem controller firmware version 07.10.xx.xx or later, use the following formulas to determine the maximum queue depth:

DS4800: $4096 / (\text{number-of-hosts} * \text{LUNs-per-host})$. For example, a DS4800 system with four hosts, each with 32 LUNs, would have a maximum queue depth of 32: $4096 / (4 * 32) = 32$.

DS4700 or DS4200: $2048 / (\text{number-of-hosts} * \text{LUNs-per-host})$. For example, a DS4700 system or a DS4200 system with four hosts, either with 32 LUNs, would have a maximum queue depth of 16: $2048 / (4 * 32) = 16$.

Formula for controller firmware versions 05.4x.xx.xx, or 06.1x.xx.xx to 06.6x.xx.xx

On DS4000 or DS5000 storage systems that are running storage subsystem controller firmware versions 05.4x.xx.xx, or 06.1x.xx.xx to 06.6x.xx.xx, use the following formula to determine the maximum queue depth: $2048 / (\text{number-of-hosts} * \text{LUNs-per-host})$. For example, a system with four hosts, each with 32 LUNs, would have a maximum queue depth of 16: $2048 / (4 * 32) = 16$.

Formula for controller firmware version 05.30.xx.xx

On DS4000 or DS5000 storage systems that are running storage subsystem controller firmware version 05.30.xx.xx or earlier, use the following formula to determine the maximum queue depth: $512 / (\text{number-of-hosts} * \text{LUNs-per-host})$. For example, a system with four hosts, each with 32 LUNs, would have a maximum queue depth of 4: $512 / (4 * 32) = 4$.

Changing the queue depth for Windows

You can use the QLogic SANsurfer program to modify the Host Adapter Settings and Advanced Adapter Settings preferences from the Windows operating-system environment. However, you must restart the servers for the changes to take effect.

Alternatively, to change the queue-depth setting for a QLogic adapter in a Microsoft Windows operating-system environment, you must select the **Configuration Settings** menu in Fast!UTIL and then select **Advanced Adapter Settings** to access the **Execution Throttle**.

Changing the queue depth for AIX

You can change the `queue_depth` attribute for AIX with the `chdev -l` command, as shown in the following example.

```
# chdev -l hdiskX -a queue_depth=y -P
```

where *X* is the name of the hdisk and *y* is the queue-depth setting.

Note: Use the `-P` flag to make the changes permanent in the Customized Devices object class.

Disabling cache mirroring

Attention: Before you disable cache mirroring, back up all data. Disabling cache mirroring might cause data loss if a controller fails, controller is reset, or is powered off.

To disable cache mirroring in Storage Manager, complete the following steps:

1. On the **Logical** or **Physical** tab of the Subsystem Management window, right-click the logical drive on which you want to disable cache mirroring, and select **Change > Cache Settings**.
2. In the Change Cache Settings window, clear the **Enable write caching with mirroring** check box.
3. Click **OK**.

Note: For AIX operating systems, when a LUN is opened that is running with write cache enabled and cache mirroring disabled, an FCP array warning message is displayed. The warning is displayed again every 24 hours until cache mirroring is enabled again.

Using dynamic capacity expansion and dynamic logical drive expansion

Dynamic logical drive expansion (DVE) increases the size of a logical drive. For you to perform a DVE, free capacity must be available in the array. If there is not, you can perform a dynamic capacity expansion (DCE) to increase the array capacity by adding drives. After you make sure that there is sufficient free capacity within the array, you can perform a DVE operation.

Performing a dynamic capacity expansion

To increase the capacity on the array, complete the following steps to perform a DCE.

Note: For more information about this procedure, see the Storage Manager online help.

1. On the **Logical** or **Physical** tab of the Subsystem Management window, right-click an array and select **Add Free Capacity (Drives)**.
2. In the Add Free Capacity (Drives) window, select one or two available drives and click **Add**.

Performing a dynamic logical drive expansion

Before you perform a Dynamic Logical Drive Expansion(DVE), make sure that there is available free capacity in the array. Check the **Logical** or **Physical** tab of the Subsystem Management window to check the amount of available free capacity. If there is not enough free capacity and extra drives are available, you can add one or more drives to the array by performing a dynamic capacity expansion (DCE) operation before you perform the DVE operation.

Note:

1. You cannot resize the logical drive while the array is activated in classic or enhanced concurrent mode.
2. You cannot resize the root array.

To increase the size of a logical drive, complete the following steps to perform a DVE.

Note: For more information about this procedure, see the Storage Manager online help.

1. From the **Logical** or **Physical** tab of the Subsystem Management window, right-click the logical drive and select **Increase Capacity**. The Increase Logical Drive Capacity – Additional Instructions window opens.
2. Read the additional instructions and click **OK**. The Increase Logical Drive Capacity window opens.
3. Type the amount by which you want to increase the logical drive, and click **OK**. A clock icon is displayed on every logical drive within the array. You must wait for the process to be completed before you can begin any host intervention.

Note: If the storage subsystem is busy, the process might take several hours.

4. Type the following commands to rescan the logical drive on the host:

```
# cd /sys/block/sdXX/device
# echo 1 > rescan
```

where XX is the device name.

5. Check the size of the logical drive, using the procedure that is described in "Checking LUN size" on page 169.
6. Remount the logical drive.

Veritas Storage Foundation with SUSE Linux Enterprise Server

Boot time is sometimes increased by LVM scanning, and Veritas Storage Foundation environment does not require an LVM scan. Therefore, the LVM scan with SLES 10 SP2 or later must be disabled. Use the following procedure to disable the LVM scan.

Note:

- In the Veritas Storage Foundation Linux environment, the default host type must be set to 13 (LNXCLVMWARE or LNXCLUSTER depending on the controller firmware version).
 - IBM supports the DMP A/P-F ASL/APM only, not the A/P-C ASL.
 - During boot, before DMP is loaded, I/O probes that go to the non-owning controller generate timeout errors. These boot-time errors are unavoidable and not significant.
1. In the file `/etc/lvm/lvm.conf`, change the line `filter = ["a./"]` to `filter = ["r|/dev/./by-path/.*|", "r|/dev/./by-id/.*|", "r|/dev/sd.*|", "a./"]`.
 2. If the root/swap is in an LVM volume, complete the following tasks:
 - Add your specific device to the filter in step 1 to make sure that the appropriate volumes are scanned.
 - After you complete step 1, run `mkinitrd` and use the new `initrd` image for future boots.

Veritas Storage Foundation 5.0 with Red Hat Enterprise Linux

The following procedure is required to enable the RDAC module on RHEL 5.3 for Storage Foundation 5.0 only. The module is integrated in Storage Foundation 5.1 and later. The `scsi_dh_RDAC` module provides the support for RDAC devices. It eliminates the time delay and some of the error messages during the boot or probe process.

Note:

1. In the Veritas Storage Foundation Linux environment, the default host type must be set to 13 (LNXCLVMWARE or LNXCLUSTER depending on the controller firmware version).
2. IBM supports the DMP A/P-F ASL/APM only, not the A/P-C ASL.
3. During boot, before DMP is loaded, I/O probes that go to the non-owning controller generate timeout errors. These boot-time errors are unavoidable and not significant.
4. The following procedure works with the IBM NVSRAM because the `scsi_dh_RDAC` module is VID/PID-dependent.

Enabling the RDAC module on RHEL 5.3 for Storage Foundation 5.0

To enable the RDAC module on RHEL 5.3 for Storage Foundation 5.0, complete the following steps:

1. Disable all of the storage subsystem storage ports so that the HBA cannot detect them.
2. Install Storage Foundation.
3. Run `mkinitrd` to include the `scsi_dh_rdac` module:

```
mkinitrd $resultant_initrd_image_file $kernel_version --preload=scsi_dh_rdac
```

For example:

```
mkinitrd /boot/my_image 2.6.18-118.el5 --preload=scsi_dh_rdac
```

Note: The `uname -r` command gives the kernel version.

4. Change the boot loader to use the new `initrd` image. For IBM i and System p servers, the `initrd` image name is `yaboot`. For System x servers, the image name is `grub`.
5. Shut down the host server.
6. Enable the storage subsystem so that the HBA recognizes the storage configuration.
7. Start the host server.

Unloading the RDAC module on RHEL 5.3 for Storage Foundation 5.0

To unload the module after the device probe and attach process, complete the following steps during the system-boot process:

1. Create a `/etc/rc3.d` script, as in the following example:

```
# vi /etc/init.d/rm_rdac
```

```
-----
## this script is used for detaching the scsi_dh_rdac module
## for each LUN
## this script has dependency on lsscsi command and this lsscsi
## should be available for this script to successfully execute.
#!/bin/bash
echo "detaching the scsi_dh_rdac module"
for i in /sys/block/sd*/device/dh_state
do
    if [[ "`cat $i`" = "rdac" ]]
    then
        echo detach > $i
    fi
done

modprobe -r scsi_dh_rdac
echo "detached successfully"
-----
```

2. Insert the script at the correct location under `/etc/rc3.d`, before the VCS VxFen Driver startup script (the VxFen Driver default start script is `/etc/rc2.d/S68vxfen`). If the system is not running VCS, insert the script after the `/etc/rc3.d/S50vxvm-recover` script.

```
# ln -s /etc/init.d/rm_rdac /etc/rc.d/rc3.d/S57rm_rdac
# ln -s /etc/init.d/rm_rdac /etc/rc.d/rc5.d/S57rm_rdac
```

Checking LUN size

To check the size of a LUN in AIX, complete the following steps:

1. Type the following commands:

```
#cd /sys/block/sdXX
# cat size
```

where *XX* is the device name. A number is displayed, as in the following example:

```
8388608
```

2. Multiply this number by 512 (bytes) to calculate the size of the LUN, as shown in the following example:

```
8388608 * 512 = 4294967296 (~ 4GB)
```

In the example, the LUN size is approximately 4 GB.

Redistributing logical drives

In a failover condition where logical drives have failed over to their secondary controller path, some configurations require a manual intervention to move these drives back after the error has been resolved. The need for this intervention depends on the host multipath driver that is installed and whether ADT (Auto Drive Transfer) is enabled. By default, ADT is disabled in AIX and Windows, but their multipath drivers can automatically recover. By default, ADT is enabled in Linux, but the MPP driver can do the same automatic recover; ADT must be disabled if you use that driver.

To redistribute logical drives manually to their preferred paths in the Subsystem Management window, click **Advanced > Recovery > Redistribute Logical Drives**.

To redistribute logical drives on AIX, HP-UX, or Solaris, complete the applicable procedure in this section.

Redistributing logical drives on AIX

If you enabled autorecovery on the AIX host, you do not have to redistribute logical drives manually after a controller failover. However, if you have a heterogeneous host environment, you might have to redistribute logical drives manually. Hosts that do not support some form of autorecovery, or AIX hosts in which autorecovery is disabled, do not automatically redirect logical drives to the preferred paths.

For troubleshooting information about disk array errors on AIX, see “Resolving disk array errors on AIX” on page 266 in Chapter 7, “Troubleshooting,” on page 245.

Complete the following steps to manually redistribute logical drives to their paths:

1. Repair or replace any faulty components. For more information, see the *Installation, User's, and Maintenance Guide* that came with the storage subsystem.
2. In the Subsystem Management window, click **Advanced > Recovery > Redistribute Logical Drives** to redistribute logical drives to their preferred paths.

Note: If a large number of LUNs are configured on the storage subsystem, the redistribution of the logical drives might take 60 minutes or longer.

3. Run the **fget_config** command to verify the active paths, as shown in the following example.

```
# fget_config -l dar0
dac0 ACTIVE dac1 ACTIVE
dac0-hdisk1
dac0-hdisk2
dac0-hdisk3
dac1-hdisk4
dac1-hdisk5
dac1-hdisk6
dac1-hdisk7
dac0-hdisk8
```

Redistributing logical drives on HP-UX

Auto Drive Transfer (ADT) is enabled, by default, on HP-UX hosts. If a failure occurs that initiates a controller failover, ADT redirects I/O to the available controller. ADT does not require manual redistribution.

Important: If a failure occurs in a heterogeneous host environment, the HP-UX host with ADT enabled automatically redistributes its LUNs when the path becomes available. However, you must redistribute logical drives manually on any host in which ADT is not enabled. Failure to do so leaves the storage subsystem in a Needs Attention state, because hosts that do not support ADT and hosts in which ADT is disabled do not automatically redirect I/O to the preferred controller. In this case, the Storage Manager Recovery Guru indicate which host platform is associated with the LUN that is in a failover state.

Note: DS5000 storage subsystems are not ALUA-compliant. DS5000 subsystems have Target Port Group Support (TPGS), which is a similar SCSI protocol that directs I/O to preferred ports. For HP-UX 11.31, the default HP-UX host type must be changed to the TPGS host type HPXTPGS.

To turn on TPGS support and change the host type, complete the following steps:

1. Change the operating-system type for the DS5000 storage subsystem from HP-UX to HPXTPGS.
2. Change the load balancing to **Default, round-robin**.
3. Verify that the changes are correct. The following example shows one of the LUNs that has the correct four active paths and four standby paths.

```
# scsimgr get_info all_lpt -D /dev/rdisk/asmlai | grep -e STATUS -e
'Open close state'
```

```
STATUS INFORMATION FOR LUN PATH : lunpath306
Open close state                = ACTIVE
STATUS INFORMATION FOR LUN PATH : lunpath344
Open close state                = STANDBY
STATUS INFORMATION FOR LUN PATH : lunpath420
Open close state                = STANDBY
STATUS INFORMATION FOR LUN PATH : lunpath326
Open close state                = ACTIVE
STATUS INFORMATION FOR LUN PATH : lunpath346
Open close state                = ACTIVE
STATUS INFORMATION FOR LUN PATH : lunpath213
Open close state                = ACTIVE
STATUS INFORMATION FOR LUN PATH : lunpath273
Open close state                = STANDBY
STATUS INFORMATION FOR LUN PATH : lunpath179
Open close state                = STANDBY
```

4. Use the SAN Fibre Channel switch-monitoring tools to verify that the I/O loads are distributed properly.

Redistributing logical drives on Solaris

If you select Solaris as the host type when you define the host and host port, Auto Drive Transfer (ADT) is disabled on Solaris hosts. In this case, if a failure occurs that initiates a controller failover, you must redistribute logical drives to their preferred paths manually.

To redistribute logical drives to their preferred paths manually, complete the following steps:

1. Repair or replace any faulty components. For more information, see the *Installation, User's and Maintenance Guide* that came with the storage subsystem.
2. In the Subsystem Management window, click **Advanced > Recovery > Redistribute Logical Drives** to redistribute logical drives to their preferred paths.

Replacing hot-swap HBAs

Attention: If you do not follow this procedure as it is documented here, data availability might be lost. You must read and understand all of the steps in this section before you begin the HBA hot-swap procedure.

This section describes the procedure for hot-swapping Fibre Channel host bus adapters (HBAs) on a System p server.

The following list provides an overview of this section:

- “Replacing hot-swap HBAs on AIX”
- “Replacing IBM HBAs on Linux” on page 175
- “Replacing a PCI hotplug HBA” on page 178
- “Mapping the new WWPN to the storage subsystem for AIX and Linux” on page 179
- “Completing the HBA hot-swap procedure” on page 179

Replacing hot-swap HBAs on AIX

Attention: Any deviations from these notes and procedures might cause a loss of data availability.

Review the following list of issues and restrictions before you perform a hot-swap operation for AIX.

- The autorecovery attribute of the dar must be set to no. Autorecovery is a dynamically set feature that can be turned back on after the hot-swap procedure is complete. Failure to disable autorecovery mode during a hot-swap procedure can cause loss of access to data.
- Do not redistribute logical drives to the preferred path until you verify that the HBA replacement succeeded and that the subsequent configuration was performed correctly. If you redistribute the logical drives before you verify that the hot swap and configuration were successful can cause a loss of access to data.
- The only supported hot-swap scenario involves the replacement of a defective HBA with the same HBA model, and in the same PCI slot. Do not insert the defective HBA into any other system, even if the HBA is found not to be defective. Always return the HBA to IBM.

Important: As of the date of this document, no other variations of replacement scenarios are supported.

- Hot swap is not supported in single-HBA configurations.

Preparing for the HBA hot swap on AIX:

To prepare for the hot swap, complete the following procedures:

Collecting system data

To collect data from the system, complete the following steps:

1. Type the following command:

```
# lsdev -C |grep fcs
```

The output is similar to the following example.

```
fcs0      Available 17-08      FC Adapter
fcs1      Available 1A-08      FC Adapter
```

2. Type the following command:

```
# lsdev -C |grep dac
```

The output is similar to the following example.

```
dac0      Available 17-08-02    1815    DS4800 Disk Array Controller
dac1      Available 1A-08-02    1815    DS4800 Disk Array Controller
```

3. Type the following command for each of the fcs devices:

```
# lscfg -vpl fcsX
```

where *X* is the number of the fcs device. The output looks similar to the following example.

```
lscfg -vpl fcs0
fcs0      U0.1-P1-I1/Q1  FC Adapter

Part Number.....09P5079
EC Level.....A
Serial Number.....1C21908D10
Manufacturer.....001C
Feature Code/Marketing ID...2765
FRU Number.....09P5080
Network Address.....10000000C92D2981
ROS Level and ID.....02C03951
Device Specific.(Z0).....2002606D
Device Specific.(Z1).....00000000
Device Specific.(Z2).....00000000
Device Specific.(Z3).....03000909
Device Specific.(Z4).....FF401210
Device Specific.(Z5).....02C03951
Device Specific.(Z6).....06433951
Device Specific.(Z7).....07433951
Device Specific.(Z8).....20000000C92D2981
Device Specific.(Z9).....CS3.91A1
Device Specific.(ZA).....C1D3.91A1
Device Specific.(ZB).....C2D3.91A1
Device Specific.(YL).....U0.1-P1-I1/Q1
```

PLATFORM SPECIFIC

```
Name: Fibre Channel
Model: LP9002
Node: Fibre Channel@1
Device Type: fcp
Physical Location: U0.1-P1-I1/Q1
```

4. Type the following command:

```
# lsdev -C |grep dar
```


The output looks similar to the following example.

```
# dar0      Available      1815    DS4800 Disk Array Router
dar1      Available      1815    DS4800 Disk Array Router
```

5. Type the following command to list the attributes of each dar found on the system:

```
# lsattr -El darX
```

where *X* is the number of the dar. The output looks similar to the following example.

```
lsattr -El dar0
act_controller  dac0,dac2  Active Controllers          False
all_controller  dac0,dac2  Available Controllers       False
held_in_reset   none       Held-in-reset controller   True
load_balancing  no         Dynamic Load Balancing     True
autorecovery    no         Autorecover after failure is corrected True
hlthchk_freq    600       Health check frequency in seconds True
aen_freq        600       Polled AEN frequency in seconds True
balance_freq    600       Dynamic Load Balancing frequency in seconds True
fast_write_ok   yes       Fast Write available       False
cache_size      1024     Cache size for both controllers False
switch_retries  5         Number of times to retry failed switches True
```

Verifying that autorecovery is disabled

Before you perform the hot swap, complete the following steps to make sure that autorecovery is disabled on every dar that is involved with the HBA that you want to hot swap:

1. Type the following command to identify all of the dacs that are involved with the HBA:

```
# lsdev -C|grep 11-08
```

The output looks similar to the following example.

```
# lsdev -C|grep 11-08
fcs0      Available 11-08      FC Adapter
fscsi0    Available 11-08-01    FC SCSI I/O Controller Protocol Device
dac0      Available 11-08-01    1742    (700) Disk Array Controller
hdisk1    Available 11-08-01    1742    (700) Disk Array Device
hdisk3    Available 11-08-01    1742    (700) Disk Array Device
hdisk5    Available 11-08-01    1742    (700) Disk Array Device
hdisk7    Available 11-08-01    1742    (700) Disk Array Device
hdisk8    Available 11-08-01    1742    (700) Disk Array Device
```

2. Check the **lsattr** command output that you collected in step 5 of the procedure “Collecting system data” on page 172. In the **lsattr** output, identify the dars that list the dacs that you identified in step 1 of this procedure.
3. Type the following command for each dar that you identified in step 2:

```
# lsattr -El darX |grep autorecovery
```

where *X* is the number of the dar. The output looks similar to the following example.

```
# lsattr -El dar0 |grep autorecovery
autorecovery  no         Autorecover after failure is corrected    True
```

4. In the **lsattr** command output, verify that the second word is no. If the second word is yes, autorecovery is currently enabled.

Important: For each dar on which autorecovery is enabled, you must disable it by setting the autorecovery ODM attribute to no. See “Using the lsattr command to view ODM attributes” on page 306 to learn how to change

attribute settings. Do not proceed with the hot-swap procedure until you complete this step and verify that autorecovery is disabled.

Replacing the hot-swap HBA:

Attention: If you do not follow this procedure as documented here, data availability might be lost. You must read and understand all of the steps in this section before you begin the HBA hot-swap procedure.

To replace the hot-swap HBA, complete the following steps:

1. Type the following command to put the HBA that you want to replace into the Defined state:

```
# rmdev -R1 fcsX
```

where *X* is the number of the HBA. The output is similar to the following example.

```
rmdev -R1 fcs0
fcnet0 Defined
dac0 Defined
fscsi0 Defined
fcs0 Defined
```

For Linux operating systems, type the following command to identify the PCI hotplug slot:

```
# drslot_chrp_pci -i -s slot-name
```

where *slot-name* is the name of the slot for the HBA that you are replacing, for example, U7879.001.DQD014E-P1-C3.

The LED at slot *slot-name* flashes, and the following message is displayed.

```
The visual indicator for the specified
PCI slot has been set to the identify
state. Press Enter to continue or
enter x to exit.
```

2. In the AIX smit menu, initiate the process that is required for the HBA hot swap by clicking **smit > Devices > PC Hot Plug Manager > Replace/Remove a PCI Hot Plug Adapter**.
3. In the Replace/Remove a PCI Hot Plug Adapter window, select the targeted HBA. A window opens and displays instructions for replacing the HBA.
4. Follow the smit instructions to replace the HBA.

Note: Do not reinstall the Fibre Channel cable at this time.

5. If the steps in this procedure are completed successfully up to this point, you get the following results:
 - The defective HBA is removed from the system.
 - The replacement FC HBA is turned on.
 - The associated fcsX device is in the Defined state.

Before you continue, verify that these results have been obtained.

6. Install the Fibre Channel loop on the replacement HBA.
7. Type the following command to put the HBA into the Active state:

```
# cfgmgr
```

Note: The new HBA is placed in the default group. If hdisks are assigned to the default group, the HBA generates a new dar and dac, which causes a split. Issue the **rmdev** command to remove the new dar and dac after you map the WWPN.

8. Type the following command to verify that the fcs device is now available:

```
# lsdev -C |grep fcs
```

9. Type the following command to verify or upgrade the firmware on the replacement HBA to the correct level:

```
# lscfg -vpl fcsX
```

where *X* is the number of the fcs.

10. Record the 16-digit number that is associated with Network Address, as it was displayed in the output of the command that you used in step 9. This network address number is used in the next procedure, “Mapping the new WWPN to the storage subsystem for AIX and Linux” on page 179.
11. Type the following command to put the HBA back into the Defined state:

```
# rmdev -Rl fcsX
```

After you complete this procedure, continue to “Mapping the new WWPN to the storage subsystem for AIX and Linux” on page 179.

Replacing IBM HBAs on Linux

This section provides requirements and procedures for replacing IBM host bus adapters in System p servers, using PCI hotplug tools.

Preparing for the IBM HBA hot swap on Linux: To prepare for the hot swap, complete the following procedures:

Verifying the PCI hotplug tools

Make sure that the following tools are installed in the `/usr/sbin` directory:

- `lsslot`
- `drslot_chrp_pci`

If these tools are not installed, complete the following steps to install them:

1. Make sure that `rdist-6.1.5-792.1` and `compat-2004.7.1-1.2` are installed from the SLES 9 media.
2. To find the PCI Hotplug Tools rpm files, go to <http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/>.
3. On the website, select the applicable link for your operating system. Download and install the following rpm files:
 - `librtas-1.3.1-0.ppc64.rpm`
 - `rpa-pci-hotplug-1.0-29.ppc64.rpm`
4. Type the following command to install each rpm file:

```
# rpm -Uvh <filename>.rpm
```

where *<filename>* is the name of the rpm file.

Verifying that the PCI core is installed

The PCI core must be installed on the system. Type the following command to verify that it is installed:

```
# ls -l /sys/bus/pci/slots
```

If the PCI core is installed, the output looks similar to the following example:

```
elm17c224:/usr/sbin # ls -l /sys/bus/pci/slots
total 0
drwxr-xr-x  8 root root 0 Sep  6 04:29 .
drwxr-xr-x  5 root root 0 Sep  6 04:29 ..
drwxr-xr-x  2 root root 0 Sep  6 04:29 0000:00:02.0
drwxr-xr-x  2 root root 0 Sep  6 04:29 0000:00:02.4
drwxr-xr-x  2 root root 0 Sep  6 04:29 0000:00:02.6
drwxr-xr-x  2 root root 0 Sep  6 04:29 0001:00:02.0
drwxr-xr-x  2 root root 0 Sep  6 04:29 0001:00:02.6
drwxr-xr-x  2 root root 0 Sep  6 04:29 control
```

If the /sys/bus/pci/slots directory does not exist, the PCI core is not installed.

Verifying that the rpaphp driver is installed

The rpaphp driver must be installed on the system. Type the following command to verify that it is installed:

```
ls -l /sys/bus/pci/slots/*
```

If the rpaphp driver is installed, the output looks similar to the following example:

```
elm17c224:/usr/sbin # ls -l /sys/bus/pci/slots/*
/sys/bus/pci/slots/0000:00:02.0:
total 0
drwxr-xr-x  2 root root  0 Sep  6 04:29 .
drwxr-xr-x  8 root root  0 Sep  6 04:29 ..
-r--r--r--  1 root root 4096 Sep  6 04:29 adapter
-rw-r--r--  1 root root 4096 Sep  6 04:29 attention
-r--r--r--  1 root root 4096 Sep  6 04:29 max_bus_speed
-r--r--r--  1 root root 4096 Sep  6 04:29 phy_location
-rw-r--r--  1 root root 4096 Sep  6 04:29 power
```

Using the lsslot tool to list slot information: Before you replace an HBA by using PCI hotplug, you can use the lsslot tool to list information about the I/O slots. This section describes how to use lsslot and provides examples. Use the lsslot tool according to the following guidelines.

Syntax for the lsslot tool

The lsslot syntax is shown in the following example.

```
lsslot [ -c slot | -c pci [ -a | -o] ] [ -s drc-name ] [ -F delimiter ]
```

Options for the lsslot tool

The lsslot options are shown in the following list:

No options

Displays all DR slots

-c slot Displays all DR slots

-c pci Displays all PCI hotplug slots

-c pci -a

Displays all available (empty) PCI hotplug slots

-c pci -o

Displays all occupied PCI hotplug slots

-F Uses delimiter to delimit columns

Listing PCI hotplug slots with the lsslot command

This section shows the command lines that you can use to list all PCI hotplug slots, all empty PCI hotplug slots, or all occupied PCI hotplug slots. You can also view detailed information about a PCI hotplug device.

Note: In the *Device(s)* columns of the command-line outputs, the PCI devices in the slots are listed in the following format: *xxxx:yy:zz.t* (for example, 0001:58:01.1).

Type the following command to list all PCI hotplug slots:

```
# lsslot -c pci -a
```

The resulting output looks similar to the following example.

# Slot	Description	Device(s)
U7879.001.DQD014E-P1-C1	PCI-X capable, 64 bit, 133MHz slot	Empty
U7879.001.DQD014E-P1-C2	PCI-X capable, 64 bit, 133MHz slot	0002:58:01.0
U7879.001.DQD014E-P1-C3	PCI-X capable, 64 bit, 133MHz slot	0001:40:01.0
U7879.001.DQD014E-P1-C4	PCI-X capable, 64 bit, 133MHz slot	Empty
U7879.001.DQD014E-P1-C5	PCI-X capable, 64 bit, 133MHz slot	Empty
U7879.001.DQD014E-P1-C6	PCI-X capable, 64 bit, 133MHz slot	0001:58:01.0

0001:58:01.1

Type the following command to list all empty PCI hotplug slots:

```
# lsslot -c pci -a
```

The resulting output looks similar to the following example.

# Slot	Description	Device(s)
U7879.001.DQD014E-P1-C1	PCI-X capable, 64 bit, 133MHz slot	Empty
U7879.001.DQD014E-P1-C4	PCI-X capable, 64 bit, 133MHz slot	Empty
U7879.001.DQD014E-P1-C5	PCI-X capable, 64 bit, 133MHz slot	Empty

Type the following command to list all occupied PCI hotplug slots:

```
# lsslot -c pci -o
```

The resulting output looks similar to the following example.

# Slot	Description	Device(s)
U7879.001.DQD014E-P1-C2	PCI-X capable, 64 bit, 133MHz slot	0002:58:01.0
U7879.001.DQD014E-P1-C3	PCI-X capable, 64 bit, 133MHz slot	0001:40:01.0
U7879.001.DQD014E-P1-C6	PCI-X capable, 64 bit, 133MHz slot	0001:58:01.0

0001:58:01.1

To see detailed information about a PCI hotplug device, complete the following steps:

1. Select a device number from the output of `# lsslot -c pci -o`, as seen in the preceding output example.
2. Type the following command to show detailed information about the device:

```
# lspci | grep xxxx:yy:zz.t
```

where *xxxx:yy:zz.t* is the number of the PCI hotplug device. The resulting output looks similar to the following example.

```
0001:40:01.0 Ethernet controller: Intel Corp. 82545EM Gigabit
Ethernet Controller (Copper) (rev 01)
```

Replacing a PCI hotplug HBA

Complete the following procedures to replace a PCI hotplug HBA by using the `drslot_chrp_pci` command.

Attention: Before you remove the HBA, you must remove the Fibre Channel cable that is attached to the HBA. The Fibre Channel cable must remain unattached for at least 5 minutes to make sure that all I/O activity is transferred to the alternate path. Failure to remove the Fibre Channel cable can cause unwanted results.

Note: In these procedures, the variable *slot-name* refers to the slot that contains the HBA that you are replacing.

1. Type the following command to identify the PCI hotplug slot:

```
# drslot_chrp_pci -i -s slot-name
```

where *slot-name* is the name of the slot for the HBA that you are replacing, for example, U7879.001.DQD014E-P1-C3.

The LED at slot *slot-name* begins flashing, and the following message is displayed.

The visual indicator for the specified PCI slot has been set to the identify state. Press Enter to continue or enter x to exit.

2. Hot unplug, or remove, the HBA from the slot:

- a. Remove the Fibre Channel cable that is connected to the HBA, and wait for failover to be completed.

- b. After failover is complete, type the following command:

```
# drslot_chrp_pci -r -s slot-name
```

The following message is displayed.

The visual indicator for the specified PCI slot has been set to the identify state. Press Enter to continue or enter x to exit.

- c. Press Enter. The following message is displayed.

The visual indicator for the specified PCI slot has been set to the action state. Remove the PCI card from the identified slot and press Enter to continue.

- d. Press Enter.

- e. Physically remove the HBA from the slot.

- f. Type the following command to verify that the slot is empty:

```
# lsslot -c pci -s slot-name
```

If the slot is empty, the resulting output looks similar to the following example.

# Slot	Description	Device(s)
U7879.001.DQD014E-P1-C3	PCI-X capable, 64 bit, 133MHz slot	Empty

3. To hot plug the HBA into the slot, complete the following steps:

- a. Type the following command:

```
# drslot_chrp_pci -a -s slot-name
```

The following message is displayed.

The visual indicator for the specified PCI slot has been set to the identify state. Press Enter to continue or enter x to exit.

- b. Press Enter. The following message is displayed.

The visual indicator for the specified PCI slot has been set to the action state. Insert the PCI card into the identified slot, connect any devices to be configured and press Enter to continue. Enter x to exit.

- c. Insert the new HBA into the slot.

- d. Type the following command to verify that the slot is no longer empty:

```
# lsslot -c pci -s slot-name
```

If the slot is not empty, the resulting output looks similar to the following example.

# Slot	Description	Device(s)
U7879.001.DQD014E-P1-C3	PCI-X capable, 64 bit, 133MHz slot	0001:40:01.0

Mapping the new WWPN to the storage subsystem for AIX and Linux

For each storage subsystem that is affected by the hot swap, complete the following steps to map the worldwide port name (WWPN) of the HBA to the storage subsystem:

1. Start the Storage Manager and open the Subsystem Management window.
2. On the **Mappings** tab of the Subsystem Management window, click **Mappings > Show All Host Port Information**. The Host Port Information window opens.
3. Find the entry in the Host Port Information window that matches the WWPN of the *defective* HBA (the HBA that you removed), and record the alias name. Then, close the Host Port Information window.
4. On the **Mappings** tab, select the alias name of the HBA host port that you just recorded.
5. Click **Mappings > Replace Host Port**. The Replace Host Port window opens.
6. In the Replace Host Port window, verify that the current HBA Host Port Identifier, which is listed at the top of the window, matches the WWPN of the HBA that you removed.
7. Type the 16-digit WWPN, without the colon (:), of the replacement HBA in the **New Identifier** field, and click **OK**.

After you complete these steps, continue to “Completing the HBA hot-swap procedure.”

Completing the HBA hot-swap procedure

To finish the HBA hot-swap procedure, complete the applicable procedure in this section for either AIX or Linux.

Completing the HBA hot-swap procedure on AIX

1. Remove the Fibre Channel loopback plug, and insert the Fibre Channel cable that was previously attached to the HBA that you removed.

Note: Skip the following step if the HBA is directly attached to the storage subsystem or if the Fibre Channel switch zoning is based on port numbers instead of WWPNS. If you do have to modify the zoning, failure to correctly do so will prevent the HBA from accessing the storage subsystem.

2. If an HBA is attached to a Fibre Channel switch and the zoning is based on WWPN, modify the zoning information to replace the WWPN of the former HBA with the WWPN of the replacement HBA.
3. Remove the Fibre Channel loopback plug, and insert the Fibre Channel cable that was previously attached to the HBA that you removed.

Note: Skip the following step if the HBA is directly attached to the storage subsystem or if the Fibre Channel switch zoning is based on port numbers instead of WWPNs. If you do have to modify the zoning, failure to correctly do so will prevent the HBA from accessing the storage subsystem.

4. If an HBA is attached to a Fibre Channel switch and the zoning is based on WWPN, modify the zoning information to replace the WWPN of the former HBA with the WWPN of the replacement HBA.
5. Run the **cfgmgr** command to enable the HBA to register its WWPN in the Fibre Channel switch.
6. Type the following commands to verify that the replaced fcsX device and its associated dacs are placed in the Available state:

```
# lsdev -C |grep fcs

lsdev -C |grep dac
```

7. Type the following command to verify that no additional dars have been created and that the expected dars are in the Available state.

Note: With MPIO, the only time you have a dac device is when the UTM LUN is assigned.

```
# lsdev -C |grep dar
```

Attention: The presence of additional dars in the lsdev output indicates a configuration problem. If this occurs, do not continue this procedure until you correct the problem. Otherwise, data availability might be lost.

8. For each dar, type the following command to verify that affected dar attributes indicate the presence of two active dacs:

```
# lsattr -El darX|grep act_controller
```

where X is the number of the dar.

The output looks similar to the following example.

```
lsattr -El dar0|grep act_controller
act_controller dac0,dac2 Active Controllers False
```

Attention: If two dacs are not reported for each affected dar, data availability might be lost. Do not continue this procedure if two dacs are not reported for each dar. Correct the problem before you continue.

9. Redistribute volumes manually to preferred paths.
10. Verify that disks stay on the preferred path with one or both of the following methods:

Using the AIX system

Run the **mpio_get_config -Av** command, and verify that drives are on the expected path.

Using Storage Manager

In the Enterprise Management window, verify that the storage subsystems are in Optimal state. If they are not in Optimal state, make sure that any drives that are part of the subsystems that are involved with the hot-swap process are not listed in the Recovery Guru.

11. If necessary, enable autorecovery of the affected dars. See Appendix D, “Viewing and setting AIX Object Data Manager (ODM) attributes,” on page 301 to learn how to change attribute settings.

The Fibre Channel HBA hot swap is now complete.

Completing the HBA hot-swap procedure on Linux

1. Remove the Fibre Channel loopback plug, and insert the Fibre Channel cable that was previously attached to the HBA that you removed.
2. If an HBA is attached to a Fibre Channel switch *and* the zoning is based on WWPN, modify the zoning information to replace the WWPN of the former HBA with the WWPN of the replacement HBA.

Note: Skip this step if the HBA is directly attached to the storage subsystem or if the Fibre Channel switch zoning is based on port numbers instead of WWPNs. If you do have to modify the zoning, failure to correctly do so will prevent the HBA from accessing the storage subsystem.

3. If RDAC is installed, type the following command to recognize the new HBA:

```
# mppBusRescan
```

The Fibre Channel HBA hot swap is now complete.

Settings for the Windows DSM and Linux RDAC

This topic applies to both the Windows and Linux Operating Systems. The configuration settings of the failover drive that is provided with the IBM Storage Manager modifies the driver.

- For Linux, the configuration settings are in the `/etc/mpp.conf` file.
- For Windows, the configuration settings are in the `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\<DSM_Driver>\Parameters` registry key, where `<DSM_Driver>` is the name of the OEM-specific driver.

The default failover driver is `mppdsm.sys`. The host must reboot for the changes to take effect. The default values listed in the following table apply to both the Windows and Linux Operating Systems. If the default values are different for Windows and Linux, it is specified. Many of these values are overridden by the failover installer for Linux or Windows.

Important: If you change these settings from their configured values, you might lose access to the storage subsystem.

Table 32. Configuration parameters of the failover driver

Parameter Name	Default Value	Description
MaxPathsPerController	4	The maximum number of paths (logical endpoints) that are supported per controller. The total number of paths to the storage subsystem is the MaxPathsPerController value multiplied by the number of controllers. The allowed values range from 0x1 (1) to 0x20 (32) for the Windows OS, and from 0x1 (1) to 0xFF (255) for Linux RDAC. For use by Customer and Technical Support representatives only.
ScanInterval	1 (Windows) 60 (Linux)	<p>The interval time, in seconds, that the failover driver checks for these conditions:</p> <ul style="list-style-type: none"> • A change in preferred ownership for a LUN • An attempt to rebalance LUNs to their preferred paths • A change in AVT enabled status or AVT disabled status <p>For Windows, the allowed values range from 0x1 to 0xFFFFFFFF and must be specified in minutes. For Linux, the allowed values range from 0x1 to 0xFFFFFFFF and must be specified in seconds. For use by Customer and Technical Support representatives only.</p>
ErrorLevel	3	<p>This setting determines which errors to log. These values are valid:</p> <ul style="list-style-type: none"> • 0 – Display all errors • 1 – Display path failover errors, controller failover errors, retryable errors, fatal errors, and recovered errors • 2 – Display path failover errors, controller failover errors, retryable errors, and fatal errors • 3 – Display path failover errors, controller failover errors, and fatal errors • 4 – Display controller failover errors, and fatal errors <p>For use by Customer and Technical Support representatives only.</p>

Table 32. Configuration parameters of the failover driver (continued)

Parameter Name	Default Value	Description
SelectionTimeoutRetryCount	0	The number of times a selection timeout is retried for an I/O request before the path fails. If another path to the same controller exists, the I/O is retried. If no other path to the same controller exists, a failover takes place. If no valid paths exist to the alternate controller, the I/O is failed. The allowed values range from 0x0 to 0xFFFFFFFF. For use by Customer and Technical Support representatives only.
CommandTimeoutRetryCount	1	The number of times a command timeout is retried for an I/O request before the path fails. If another path to the same controller exists, the I/O is retried. If another path to the same controller does not exist, a failover takes place. If no valid paths exist to the alternate controller, the I/O is failed. The allowed values range from 0x0 to 0xa (10) for Windows, and from 0x0 to 0xFFFFFFFF for Linux RDAC. For use by Customer and Technical Support representatives only.
UaRetryCount	10	The number of times a Unit Attention (UA) status from a LUN is retried. This parameter does not apply to UA conditions due to Quiescence In Progress. The allowed values range from 0x0 to 0x64 (100) for the Windows OS, and from 0x0 to 0xFFFFFFFF For Linux RDAC. For use by Customer and Technical Support representatives only.
SynchTimeout	120	The timeout, in seconds, for synchronous I/O requests that are generated internally by the failover driver. Examples of internal requests include those related to rebalancing, path validation, and issuing of failover commands. The allowed values range from 0x1 to 0xFFFFFFFF. For use by Customer and Technical Support representatives only.

Table 32. Configuration parameters of the failover driver (continued)

Parameter Name	Default Value	Description
DisableLunRebalance	0	<p>This parameter provides control over the LUN failback behavior of rebalancing LUNs to their preferred paths. These values are possible:</p> <ul style="list-style-type: none"> • 0 – LUN rebalance is enabled for both AVT and non-AVT modes. • 1 – LUN rebalance is disabled for AVT mode and enabled for non-AVT mode. • 2 – LUN rebalance is enabled for AVT mode and disabled for non-AVT mode. • 3 – LUN rebalance is disabled for both AVT and non-AVT modes. • 4 – The selective LUN Transfer feature is enabled if AVT mode is off and ClassicModeFailover is set to LUN level 1.
S2ToS3Key	Unique key	<p>This value is the SCSI-3 reservation key generated during failover driver installation. For use by Customer and Technical Support representatives only.</p>
LoadBalancePolicy	1	<p>This parameter determines the load-balancing policy used by all logical drives managed by the Windows DSM and Linux RDAC failover drivers. These values are valid:</p> <ul style="list-style-type: none"> • 0 – Round robin with subset. • 1 – Least queue depth with subset. • 2 – Least path weight with subset (Windows only)
ClassicModeFailover	0	<p>This parameter provides control over how the DSM handles failover situations. These values are valid:</p> <ul style="list-style-type: none"> • 0 – Perform controller-level failover (all LUNs are moved to the alternate controller). • 1 – Perform LUN-level failover (only the LUNs indicating errors are transferred to the alternate controller).

Table 32. Configuration parameters of the failover driver (continued)

Parameter Name	Default Value	Description
SelectiveTransferMaxTransferAttempts	3	This parameter sets the maximum number of times that a host transfers the ownership of a LUN to the alternate controller when the Selective LUN Transfer mode is enabled. This setting prevents multiple hosts from continually transferring LUNs between controllers.
SelectiveTransferMinIOWaitTime	5	This parameter sets the minimum wait time (in seconds) that the DSM waits before it transfers a LUN to the alternate controller when the Selective LUN Transfer mode is enabled. This parameter is used to limit excessive LUN transfers due to intermittent link errors.

Wait Time Settings

When the failover driver receives an I/O request for the first time, the failover driver logs timestamp information for the request. If a request returns an error and the failover driver decides to retry the request, the current time is compared with the original timestamp information. Depending on the error and the amount of time that has elapsed, the request is retried to the current owning controller for the LUN, or a failover is performed and the request sent to the alternate controller. This process is known as a wait time. If the NotReadyWaitTime value, the BusyWaitTime value, and the QuiescenceWaitTime value are greater than the ControllerIOWaitTime value, they will have no effect.

Table 33. Parameters for Wait time settings

Parameter Name	Default Value	Description
NotReadyWaitTime	300 (Windows), 270 (Linux)	The time, in seconds, a Not Ready condition (SK 0x06, ASC/ASCQ 0x04/0x01) is allowed before failover is performed. Valid values range from 0x1 to 0xFFFFFFFF.
BusyWaitTime	600 (Windows) 270 (Linux)	The time, in seconds, a Busy condition is allowed before a failover is performed. Valid values range from 0x1 to 0xFFFFFFFF.
QuiescenceWaitTime	600 (Windows) 270 (Linux)	The time, in seconds, a Quiescence condition is allowed before a failover is performed. Valid values range from 0x1 to 0xFFFFFFFF.

Table 33. Parameters for Wait time settings (continued)

Parameter Name	Default Value	Description
ControllerIoWaitTime	600 (Windows) 120 (Linux)	Provides an upper-bound limit, in seconds, that an I/O is retried on a controller regardless of the retry status before a failover is performed. If the limit is exceeded on the alternate controller, the I/O is again attempted on the original controller. This process continues until the value of the ArrayIoWaitTime limit is reached. Valid values range from 0x1 to 0xFFFFFFFF.
ArrayIoWaitTime	600	Provides an upper-bound limit, in seconds, that an I/O is retried to the storage subsystem regardless of to which controller the request is attempted. After this limit is exceeded, the I/O is returned with a failure status. Valid values range from 0x1 to 0xFFFFFFFF.

Configuration Settings for Path Congestion Detection and Online/Offline Path States

The following configuration settings are applied using the utility `dsmUtil -o` option parameter.

Table 34. Configuration Settings for the Path Congestion Detection

Parameter	Default Value	Description
CongestionDetectionEnabled	0x0	A Boolean value that indicates whether the path congestion detection is enabled. If this parameter is not defined or is set to 0x0, the value is false, the path congestion feature is disabled, and all of the other parameters are ignored. If set to 0x1, the path congestion feature is enabled. Valid values are 0x0 or 0x1.

Table 34. Configuration Settings for the Path Congestion Detection (continued)

Parameter	Default Value	Description
CongestionResponseTime	0x0	If CongestionIoCount is 0x0 or not defined, this parameter represents an average response time in seconds allowed for an I/O request. If the value of the CongestionIoCount parameter is nonzero, then this parameter is the absolute time allowed for an I/O request. Valid values range from 0x1 to 0x10000 (approximately 18 hours).
CongestionIoCount	0x0	The number of I/O requests that have exceeded the value of the CongestionResponseTime parameter within the value of the CongestionTimeFrame parameter. Valid values range from 0x0 to 0x10000 (approximately 4000 requests).
CongestionTimeFrame	0x0	A sliding window that defines the time period that is evaluated in seconds. If this parameter is not defined or is set to 0x0, the path congestion feature is disabled because no time frame has been defined. Valid values range from 0x1 to 0x1C20 (approximately two hours).
CongestionSamplingInterval	0x0	The number of I/O requests that must be sent to a path before the nth request is used in the average response time calculation. For example, if this parameter is set to 100, every 100th request sent to a path is used in the average response time calculation. If this parameter is set to 0x0 or not defined, the path congestion feature is disabled for performance reasons—every I/O request would incur a calculation. Valid values range from 0x1 to 0xFFFFFFFF (approximately 4 billion requests).

Table 34. Configuration Settings for the Path Congestion Detection (continued)

Parameter	Default Value	Description
CongestionMinPopulationSize	0x0	The number of sampled I/O requests that must be collected before the average response time is calculated. Valid values range from 0x1 to 0xFFFFFFFF (approximately 4 billion requests).
CongestionTakeLastPathOffline	0x0	A Boolean value that indicates whether the DSM driver takes the last path available to the storage subsystem offline if the congestion thresholds have been exceeded. If this parameter is not defined or is set to 0x0, the value is false. Valid values are 0x0 or 0x1. Note: Setting a path offline with the dsmUtil utility succeeds regardless of the setting of this value.

Example Configuration Settings for the Path Congestion Detection Feature

Here's a little example section in a concept.

Note: Before path congestion detection can be enabled, you must set the CongestionResponseTime, CongestionTimeFrame, and CongestionSamplingInterval parameters to valid values.

To set the path congestion IO response time to 10 seconds: `dsmUtil -o CongestionResponseTime=10,SaveSettings`

To set the path congestion sampling interval to one minute: `dsmUtil -o CongestionSamplingInterval=60`

To enable path congestion detection: `dsmUtil -o CongestionDetectionEnabled=0x1,SaveSettings`

To use the dsmUtil -o command to set a path to Admin Offline: `dsmUtil -o SetPathOffline=0x77070001`

To use the dsmUtil -o command to set a path to online: `dsmUtil -o SetPathOnline=0x77070001`

Note: The path ID (in this example 0x77070001) is found using the dsmUtil -g command.

Setting up details on the DS5000 storage system and AIX host to support T10PI

You must set up your DS5000 storage system and AIX host to support T10PI functionality from the AIX host to the drives in the DS5000 storage subsystem using these steps.

Set up the DS5K storage box

1. Upgrade the DS5000 storage system with the firmware that supports T10PI.
2. Create and export logical unit numbers (LUNs) that support T10PI.

Note: To perform this step, you must have disks that support T10PI. These disks are generally pre-initialized to 520-byte blocks. See “T10PI capable drive attributes” on page 60 for more information.

Set up the AIX host

Note: You should use SSIC to check for supported HBAs, HBA driver and firmware version, and your AIX operating version, along with any applicable patches.

1. Install the AIX 61 TL6 SP5 or AIX 71 TL0 SP3. The AIX machine should have at least one 8Gb PCIe FC adapter (feature code 5735 or 5273) with the latest firmware that supports T10PI. It should at least have version df1000f114108a03.200305. The AIX diag utility can be used to download 8Gb PCIe FC adapter (feature code 5735 or 5273) firmware.
 - The 2 port 8Gb PCIe Fibre Channel Adapted must be used.
 - 200307 level of firmware or higher is required.
 - Feature code 5735 or low profile feature code 5273 required.
 - No support on Power Blades.
 - To download the adapter firmware, use the **diag -T download -d fcs#** command, where # is the fcs device number
 - To find the current firmware version on the Coho adapter, use the **lsmcode -cd fcs#** or **lscfg -vl fcs#** commands
 2. Enable protection on the 8Gb PCIe FC adapter (feature code 5735 or 5273) (DIF_enabled attribute of fcs device).
 - To enable protection on FC adapter, use the **chdev -l fcs# -a DIF_enabled=yes** command can be used.
 - You can also use the **smit fcsa** command can also be used to enable/disable protection on fcs#
- Note:** To disable protection on FC adapter, use the **chdev -l fcs# -a DIF_enabled=no**.
3. Ensure that you update the firmware as specified in step 1 and perform step 2 on all the 8Gb PCIe FC adapters (feature code 5735 or 5273) from which the disk has paths.
 - a. Use the **lspath -l hdisk#** command to find the FC adapters in the paths of a disk. This will show the fcs# devices (FC protocol devices).
 - b. Use the **lsdev -l fcs# -F'name parent'** command to find the parent of an fcs# device.
 - c. Enable protection.

4. Use the **chdev -l hdisk# -a DIF_protection=yes** command to enable T10 protection on a disk. The disk must support "Type 1" T10 protection.

Note: You can also use the **smit disk** command to enable/disable protection on hdisk#.

Note: You can use the **chdev -l hdisk# -a DIF_protection=no** command to disable T10 protection on a disk.

5. Use the **lsattr -El hdisk#** command to check the current value of this attribute after enabling protection. If at least one path does not support protection, then protection can't be enabled on the disk. If this attribute has a values of "unsupported," it means that:
 - some or all paths to disk won't support protection OR
 - the disk does not support protection

hdisk2 has three paths. These three paths are from fcs0, fcs2 and fcs3. You want to enable protection on these adapters. To do this:

1. Upgrade the firmware on all the fcs devices mentioned above. All of them must be 8Gb PCIe FC adapters (feature code 5735 or 5273).
2. Un-configure the child devices (fscsi0, fscsi2 and fscsi3).
3. Enable protection on fcs0, fcs2 and fcs3 adapters using the **chdev** command (**chdev -l fcs0 -a DIF_enabled=yes**).
4. Run **cfgmgr** so that all the devices will come into available state.
5. Use **chdev** command on hdisk2 to enable or disable protection (**chdev -l hdisk2 -a DIF_protection=yes**). If the disk supports protection and all paths support protection, the attribute value will be set to "yes". Otherwise, it will be set to "unsupported".

Note: If the attribute value is set to "unsupported", check all paths (all fcs adapter attributes) and check to see if the protection is enabled on the LUN when it is created on the DS5000 storage. In some cases, the attribute on the fcs adapter may show "yes" but it may not be supported due to an old 8Gb PCIe FC adapter (feature code 5735 or 5273) firmware that does not support T10 protection (BlockGuard feature).

Chapter 6. Working with full disk encryption

This chapter describes the capabilities and advantages of full disk encryption (FDE) disk drives and how to implement security on FDE-compatible storage subsystems that are equipped with FDE disks.

In addition to the information in this chapter, the *IBM Full Disk Encryption Best Practices* document describes best practices for maintaining security on a storage subsystem that is equipped with FDE drives. To access this document, go to <http://www-947.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5081492&brandind=5000028>, or complete the following steps:

1. Go to the IBM Support Portal at <http://www.ibm.com/support/entry/portal>.
2. In the **Search within all of support & downloads** field at the bottom of the webpage, type FDE and press Enter.
3. In the list of search results, click the **IBM Full Disk Encryption Best Practices - IBM System Storage** link.
4. Click the link to the PDF file to open or download the *IBM Full Disk Encryption Best Practices* document.

Note: You can also secure a disk pool, if it has FDE disks. Refer to “Securing a RAID array” on page 224. The procedure is exactly the same.

Note: The screenshots in this section are only illustrative and may differ from the actual UI depending on the Storage Manager and controller firmware versions. The following topics are addressed in this chapter:

- “Full disk encryption” on page 192
 1. “Securing data against a breach” on page 192
 2. “Choosing local or external security key management” on page 193
 3. “Before you begin” on page 210
 4. “Using security keys” on page 194
 5. “Using secure erase” on page 206
 6. “FDE security authorizations” on page 207
 7. “FDE terminology” on page 209
- “Installing and configuring the DS TKLM Proxy Code server” on page 210
 1. “Modifying the DS TKLM Proxy Code server configuration file” on page 212
 2. “Installing the DS TKLM Proxy Code” on page 215
- “Configuring disk encryption with FDE drives” on page 216
 1. “Installing FDE drives” on page 217
 2. “Enabling premium features” on page 217
 3. “Securing a RAID array” on page 224
 4. “Unlocking disk drives” on page 229
 5. “Migrating storage subsystems (head-swap) with FDE drives” on page 232
 6. “Erasing disk drives” on page 235
 7. “Global hot-spare disk drives” on page 238
 8. “Log files” on page 239
- “Frequently asked questions” on page 239

Note: Not all IBM DS storage subsystems support FDE. See the documentation that came with your storage subsystem for information about FDE compatibility.

Full disk encryption

The information in this section provides an overview of how FDE works. Subsequent sections in this chapter describe how to configure disk encryption, using internal security key management, and using external security key management.

The use of full disk encryption (FDE) secures data against threats when an FDE drive is out of its owner's control. FDE drives do not protect data from threats that occur within the data center or on the network. If an attacker gains access to a server and can access an unlocked drive, the attacker can read the clear text that comes from the drive. Remember that drive-level encryption technology does not replace the access controls of the data center; rather, it complements them.

Full disk encryption (FDE) disk drives enable you to reduce the security vulnerabilities of stored data. FDE disk drives that adhere to the Trusted Storage Group (TCG) enterprise security subsystem class specification are National Security Agency qualified and provide security with government-grade encryption.

Note: No single security implementation can effectively secure all levels of data from all threats.

Different technologies are required to protect data that is stored on hard disk drives from different threats. FDE drives protect the security of stored data through the following methods:

Securing data against a breach

If an unauthorized user gains possession of a disk drive that contains encrypted data, or the drive is removed from the data center or power is turned off, the data is protected.

Using secure erase

Secure erase provides fast, permanent erasure of data on drives that are planned for reuse or disposal.

Securing data against a breach

Drives with full disk encryption technology are security capable. Each FDE drive comes from the factory in Security Capable (security disabled) state. In this state, FDE drives behave exactly as non-FDE drives. The data that is stored on them is not protected when the drives are removed from the storage subsystem. You can move them from one storage subsystem to another without unlocking them with a security key file. They can also be used as part of a RAID array that is composed of non-encrypting (non-FDE) disks. However, a RAID array that is composed of Security Capable FDE and non-FDE drives cannot be converted into a secured RAID array at a later time, leaving the data on the FDE drives unprotected if they are removed from the storage subsystem.

The IBM storage subsystem controllers can apply security to every FDE drive in a RAID array that is composed entirely of FDE drives. Depending on the security key management method that you use (local or external), the controller firmware either creates a security key or obtains a security key from the external key manager, such as the IBM Tivoli Key Lifecycle Manager software. After the

firmware has the security key, it activates the encryption function of the drive, which causes each FDE disk drive to generate a random encryption key that is embedded on the disk.

When security is enabled, the FDE drive automatically executes full disk encryption for write and read operations. When a write operation is performed, clear text enters the disk and is encrypted before it is written to the media, using the disk encryption key. When a read operation is performed, encrypted data that is read from the media is decrypted before it leaves the drive.

During normal operation, whether the FDE drive is in Security Capable or Security Enabled state, it behaves the same as a non-encrypting disk to the storage subsystem. A security-enabled FDE drive is constantly encrypting data. Disk encryption cannot be accidentally turned off. The disk encryption key is generated by the drive itself, is stored on the disk, never leaves the disk, and is unique to that drive alone. To ensure that security is never compromised, an encrypted version of the encryption key is stored only on the disk drive. Because the disk encryption key never leaves the disk, you might not have to periodically change the encryption key, the way a user might periodically change the operating-system password.

Choosing local or external security key management

There are two methods for managing the security key for your storage subsystem: local and external security key management.

Local security key management

With local security key management, the security key is created and contained in the storage subsystem controller. Local security key management does not require additional software. To move secured drives from one storage subsystem to another, you must use the saved security key file from the original storage subsystem to unlock the drives.

To enable local security key management, complete the following tasks:

1. Follow the FDE premium feature web-activation instructions.
2. Use Storage Manager to command the storage subsystem controller to create the security key.

External security key management

Instead of using a security key created by the storage subsystem controller, external security key management uses a central key location on your network to manage keys for different storage subsystems. External security key management is facilitated by external key license manager software, such as IBM Tivoli Key Lifecycle Manager (TKLM). If you do not already have this software, you must purchase it, install it, and configure the proxy server.

With external security key management, the controllers obtain the security key from the external security key management source. This key is then obfuscated in the controller volatile memory for future use, as long as the storage subsystem power is turned on. This key is erased from volatile memory when the storage subsystem power is turned off. Because the key is not stored in the storage subsystem, the storage subsystem must have a non-FDE drive in the configuration to boot successfully; it then requests the security key from the external key management server to unlock the FDE drives.

This method provides a common and consistent key management interface; the external key license manager software also manages security keys for other storage hardware, such as secured tape drives. You do not have to access a saved security key file to move secured drives from one storage subsystem to a second storage subsystem. Rather, the external key license manager software supplies the security key that unlocks the drives automatically, if the second storage subsystem is connected to the key license manager when the drives are inserted.

To enable external security key management, complete the following tasks:

1. Install and configure the external key license manager software. See the documentation that came with the software for more information.
2. Install and configure the DS TKLM Proxy Code.
3. Configure the external key management software to receive an external key request.
4. Use Storage Manager to command the storage subsystem controller to request the security key from the external key license manager, instead of generating a local security key.
5. Configure the external key license manager software to accept an external key request.

Important:

1. Tivoli Key Lifecycle Manager is the only external security key management software that is supported on IBM DS storage subsystems.
2. Make sure that at least one non-FDE drive is installed in the storage subsystem when you use external security key management. Otherwise, if the storage subsystem power is turned off and then on again, the storage subsystem might require that you supply the security key from the saved file manually to unlock the secured FDE drives and complete the boot process.

Using security keys

With full disk encryption, the process of securing a drive consists of enabling security on the storage subsystem and then securing the specific security-capable RAID arrays where the data is stored.

The process for generating a security key depends on the type of security key management method that you use. Enabling security on the storage subsystem is a one-time process, unless the you decide at a later date to change the security key or change the method of key management. Separate security keys are not required for each individual drive, even though each FDE drive has its own unique encryption key. To enable security on the storage subsystem, you must purchase FDE drive options and an IBM DS Disk Encryption premium feature key and enable the feature in the storage subsystem, using the instructions that come with the premium feature key entitlement kit.

After the security key is created by the controllers, or is obtained from the external key management software, an encrypted version of the security key is obfuscated in the storage subsystem and cannot be viewed directly.

After you create the security key in the storage subsystem, you are asked to save the encrypted version of the security key in a backup security key file at a location that you specify. Make sure that you protect the security key file and its associated pass phrase. In addition to the saved location that you specify, the storage manager also saves a copy of the file in the default location `...\IBM_DS\client\data\`

securityLockKey in a Microsoft Windows environment or in /var/opt/SM/securityLockkey in AIX, Linux, Solaris, and HP-UX environments.

With the local security key management method, you are prompted for the security key identifier and pass-phrase. This security key identifier is appended with the storage subsystem worldwide identifier to help you identify the storage subsystem to which the security key is associated. With the external security key management method, you are prompted only for the pass-phrase. The controller uses the storage subsystem world-wide identifier to identify the storage subsystem to which the security key file is associated.

The security key file contains the encrypted security key and the security key identifier. You must provide the pass phrase during the save security key operation. The pass phrase is not stored anywhere in the storage subsystem or in the security key file. The controller uses the pass phrase to encrypt the security key before it exports the security key to the security key file. The security key identifier is stored in the security key file so that you can identify the storage subsystem to which the security key file is associated. Make sure that you protect the security key file and its associated pass phrase, because these two pieces of information can be used to unlock secured FDE drives.

To decrypt the security key in the security key file, you must provide the same pass phrase that was entered when the security key file was generated. The drive then determines whether its security key and the security key that was provided by the storage subsystem are the same. If they are the same, data can be read from and written to the security-enabled FDE drives.

Attention: The pass phrase is used only to protect the security key in the security key file. Anyone who can access the Subsystem Management window can save a copy of the security key file with a new pass phrase. Set a storage subsystem password for each storage subsystem that requires you to provide a password when any configuration changes are made, including creating and changing the security key. See “Setting a storage subsystem management password” on page 36 for instructions for setting the storage subsystem password.

If you use local security key management, the security key file provides protection against a corrupted security key or the failure of both controllers in the storage subsystem. The security key file is also needed to unlock security-enabled FDE drives when they are moved from one storage subsystem to another. In these cases, the security-enabled FDE drives remain locked until the drives are unlocked by the security key that is stored in the security key file. To decrypt the security key in the security key file, you must provide the same pass phrase that was entered when the security key file was generated. The drive then determines whether its security key and the security key that was provided by the storage subsystem are the same. If they are the same, data can be read from and written to the security-enabled FDE drives.

If you use external security key management, the security key file provides protection in the following situations:

1. If communication is lost to either the proxy server or the external key license servers when the controller unlocks the secured FDE drives
2. If the secured FDE drives are moved to or from a storage subsystem that is not managed by the same external key license manager

3. If drives must be unlocked after the power cycle of a storage subsystem configuration that has only secured FDE drives and no unsecured FDE or non-FDE drives in the configuration

After the storage subsystem controller creates the security key, the RAID arrays can be changed from a state of Security Capable to a state of Security Enabled. The Security Enabled state requires the RAID array FDE drives to be unlocked after power to the drive is turned on using the security key to access the data that is stored on the drives. Whenever power is applied to the drives in a RAID array, the drives are all placed in Security Locked state. They are unlocked only during drive initialization with the storage subsystem security key. The Security Unlocked state makes the drives accessible for the read and write activities. After they are unlocked, the drives remain unlocked until the power is removed from the drives, the drives are removed and reinserted in the drive bays, or the storage subsystem power is cycled.

After a drive is secured, the drive becomes locked if power is turned off or if it is removed. The encryption key within that drive will not encrypt or decrypt data, making the drive unreadable until it is unlocked by the controllers.

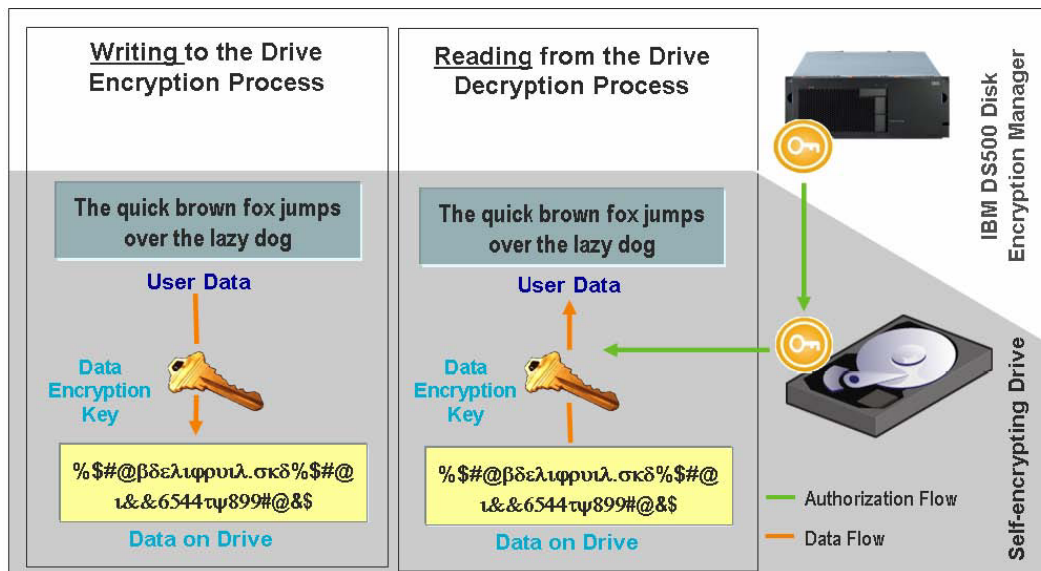


Figure 25. Security-enabled FDE drives: With the correct authorizations in place, the reading and writing of data occurs in Unlocked state

After authentications are established and security is enabled on a storage subsystem, the encryption of write operations and decryption of read operations that takes place inside the FDE drive are not apparent to the user or to the DS5000 storage subsystem controllers. However, if a secured drive is lost, removed, or stolen, the drive becomes locked, and the data that is stored on the disk remains encrypted and unreadable. Because an unauthorized user does not have the security key file and pass phrase, gaining access to the stored data is impossible.




Figure 26. A security-enabled FDE drive is removed from the storage subsystem: Without correct authorizations, a stolen FDE disk cannot be unlocked, and the data remains encrypted

Changing a security key for local security key management

The process for changing a security key depends on the type of security key management method that you use. The information in this section describes changing a security key in a local security key management configuration.

When you change a security key, a new security key is generated by the storage subsystem controller firmware. The new security key is obfuscated in the storage subsystem, and you cannot see the security key directly. The new security key replaces the previous key that is used to unlock the security-enabled FDE drives in the storage subsystem. The controller negotiates with all of the security-enabled FDE drives for the new key.

A backup copy of the security key file is always generated when you change a security key, and must be stored on some other storage medium in case of controller failure, or for transfer to another storage subsystem. You participate in creation of the security key identifier, the pass phrase, and the security key file name and location when you change the security key. The pass phrase is not stored anywhere in the storage subsystem or in the security file. The controller uses the pass phrase to encrypt the security key before it exports the security key to the security key file.

Changing a security key for external security key management

The information in this section describes changing a security key in an external security key management configuration.

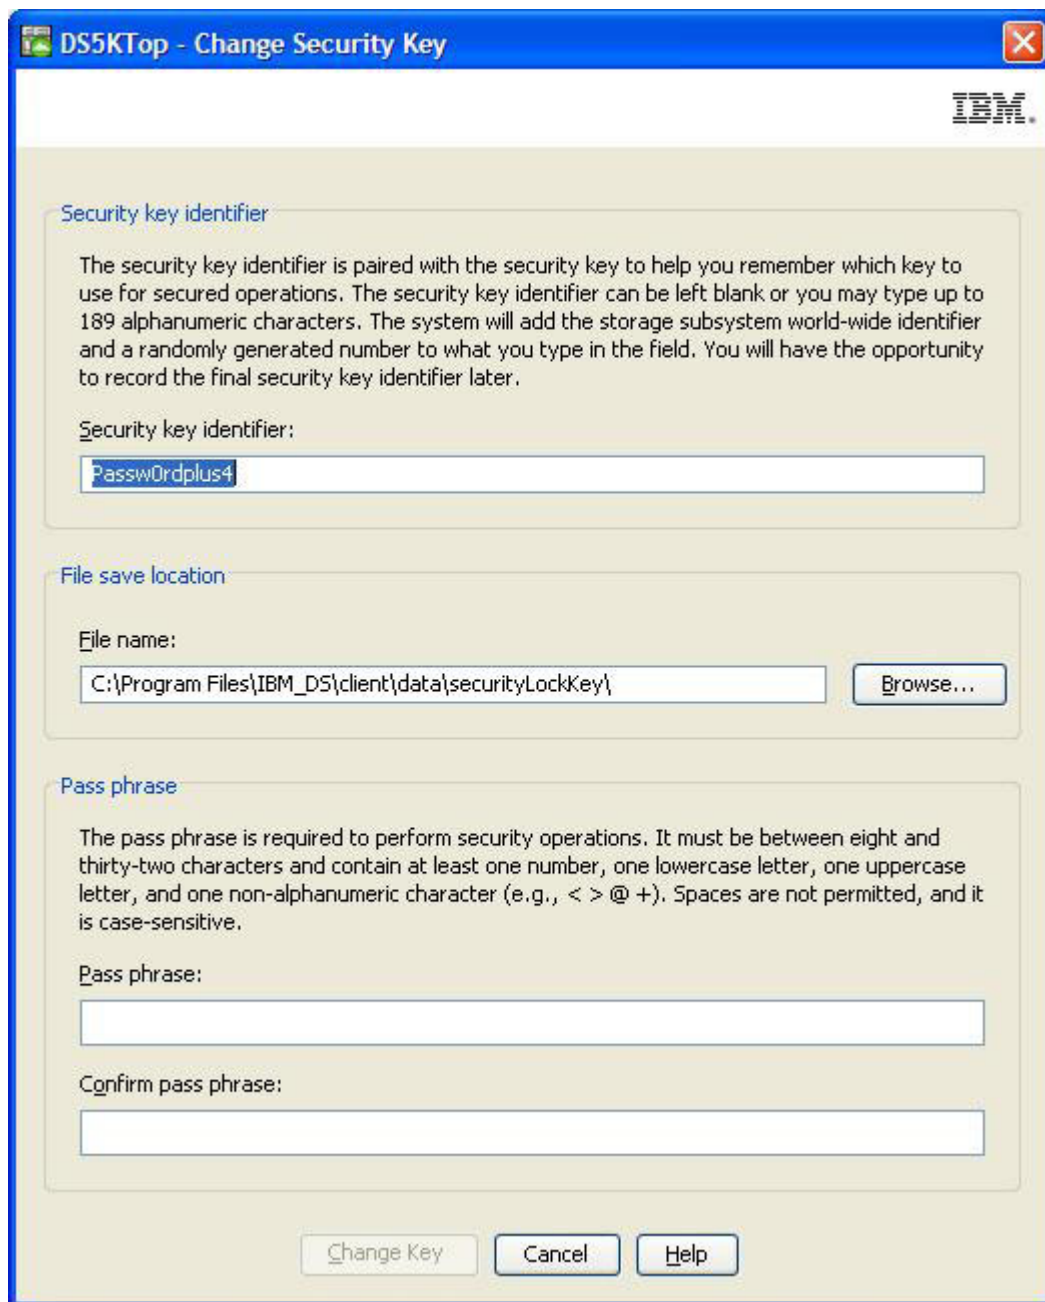
When you change the security key, the storage subsystem controller contacts the external key license manager for a new security key. Then, it negotiates the new security key with the security-enabled FDE drives. The new key is not obfuscated inside the controller. You are prompted to save the key in a security key file. The pass phrase and the security key file name and location are required to back up the security key. The pass phrase is not stored anywhere in the storage subsystem or in the security key file. The controller uses the pass phrase to encrypt the security key before it exports the security key to the security key file.

Identifying a security key file for a storage subsystem

For additional protection, the security key that is used to unlock FDE drives is not visible to the user. The security key identifier helps you identify which security key file is associated with each storage subsystem. With the local security key management method, you can provide a value of up to 189 alphanumeric characters. This value is linked with the storage subsystem worldwide identifier and a random number to form the security key identifier. In the external security key management method, you are not asked to provide a value that is used as part of the security key identifier. You can view the security key identifier during operations that involve the drive security key file, such as creating or changing the security key.

Figure 27 on page 199 shows an example of the security key identifier field when you are performing a change security key operation.

Note: With external security key management, the security key identifier cannot be modified by the user as it can with local security key management.



The image shows a Windows-style dialog box titled "DS5KTop - Change Security Key". It features the IBM logo in the top right corner. The dialog is divided into three main sections: "Security key identifier", "File save location", and "Pass phrase".

Security key identifier

The security key identifier is paired with the security key to help you remember which key to use for secured operations. The security key identifier can be left blank or you may type up to 189 alphanumeric characters. The system will add the storage subsystem world-wide identifier and a randomly generated number to what you type in the field. You will have the opportunity to record the final security key identifier later.

Security key identifier:

File save location

File name:

Pass phrase

The pass phrase is required to perform security operations. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., < > @ +). Spaces are not permitted, and it is case-sensitive.

Pass phrase:

Confirm pass phrase:

At the bottom of the dialog are three buttons: "Change Key", "Cancel", and "Help".

Figure 27. Changing the security key

The Change Security Key Complete window shows that the security key identifier that was written to the security key file has a random number appended to the security key identifier you entered in Figure 27 and the storage subsystem worldwide identifier. Figure 28 on page 200 shows an example of the random number part of the security key identifier.



Figure 28. Changing the security key - Complete

The **Security key identifier** field in the FDE Drive Properties window includes a random number that is generated by the controller when you create or change the security key. Figure 29 on page 201 shows an example of the random number. The random number is currently prefixed with 27000000. If all of the secured FDE drives in the storage subsystem have the same value in the security key identifier field, they can be unlocked by the same security key identifier.

Note: The **Security Capable** and **Secure** fields in the Drive Properties window show whether the drive is secure capable and whether it is in Secure (Yes) or Unsecured (No) state.



Figure 29. Drive properties - Secure FDE drive

Figure 30 on page 203 shows an example of the security key identifier that is displayed in the **File information** field when you select a security key back up file to unlock the secured drives in the storage subsystem. The security key identifier

or LockKeyID, shown in the file information field, contains the characters that you entered in the security key identifier field when you created or changed the security key along with the storage subsystem worldwide identifier and the randomly-generated number that appears in the security key identifier of all secured FDE drives. This information is delimited by a colon (:). For example:

Passw0rdplus3:600a0b800029ece6000000004a2d0880:600a0b800029ed8a00001aef4a2e4a73

A LockKeyID contains the following information:

- The security key identifier that you specified, for example Passw0rdplus3

Note: With external security key management, the security key identifier cannot be modified by the user as it can with local security key management. Therefore, this information will not be shown.

- The storage subsystem worldwide identifier, for example
600a0b800029ece6000000004a2d0880
- A randomly-generated number 600a0b800029ed8a00001aef4a2e4a73

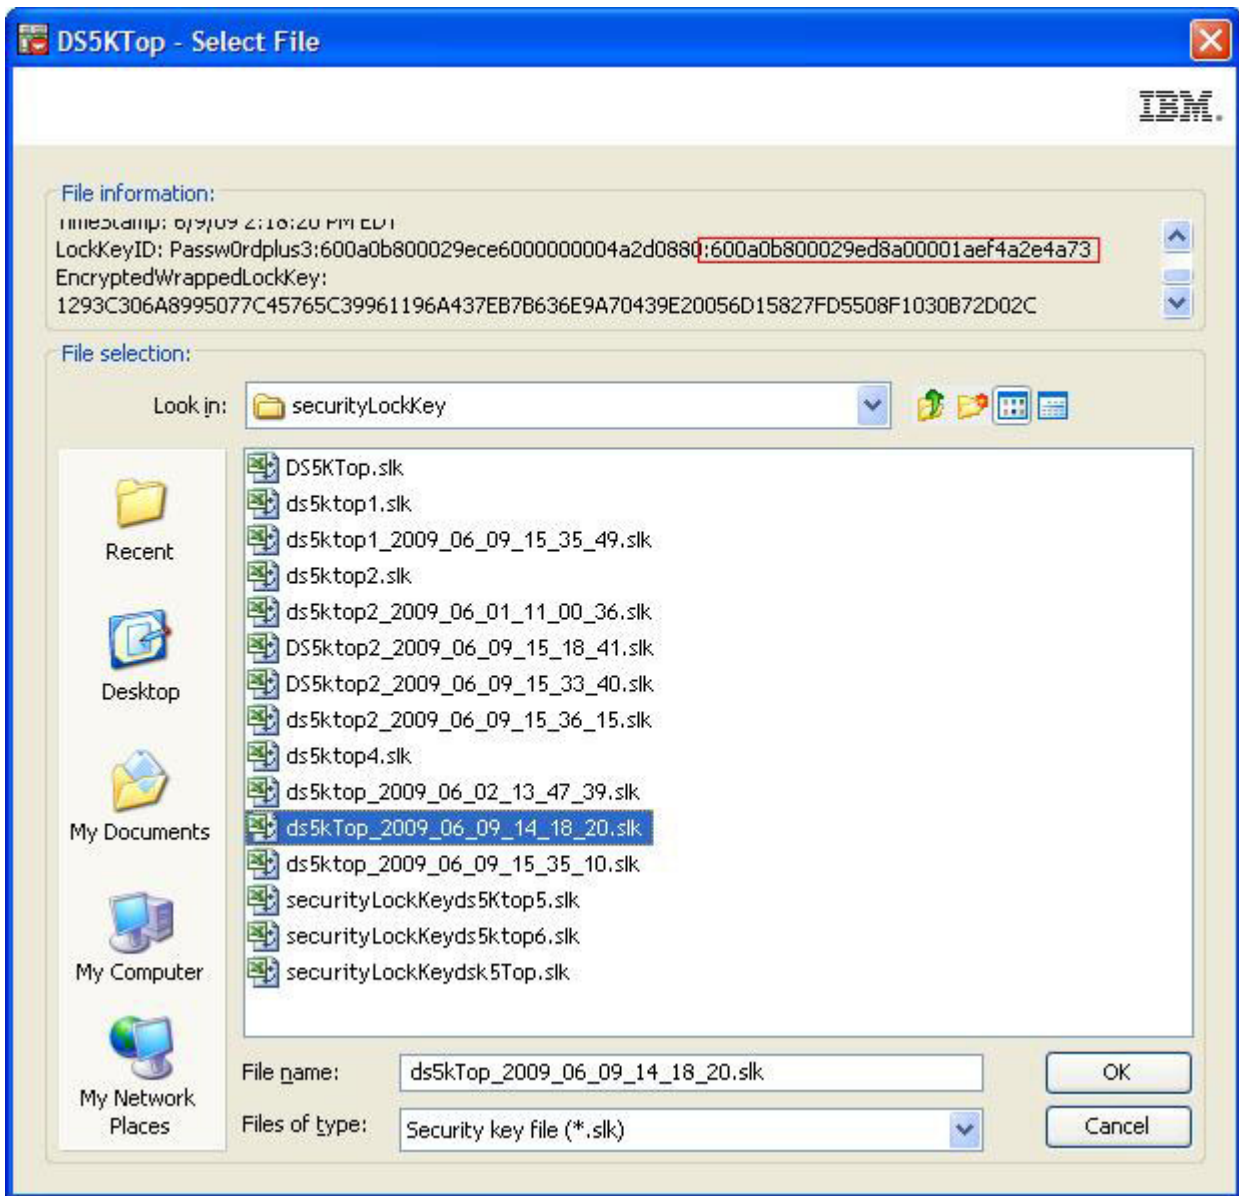


Figure 30. Select file - LockKeyID

Figure 31 on page 204 shows an example of the drive properties for an unsecured FDE drive. Note that the security key identifier field for an unsecured FDE drive is populated with zeros. Note also that the **Security Capable** field value is yes and the **Secure** field value is no, indicating that this is a security capable but unsecured FDE drive.

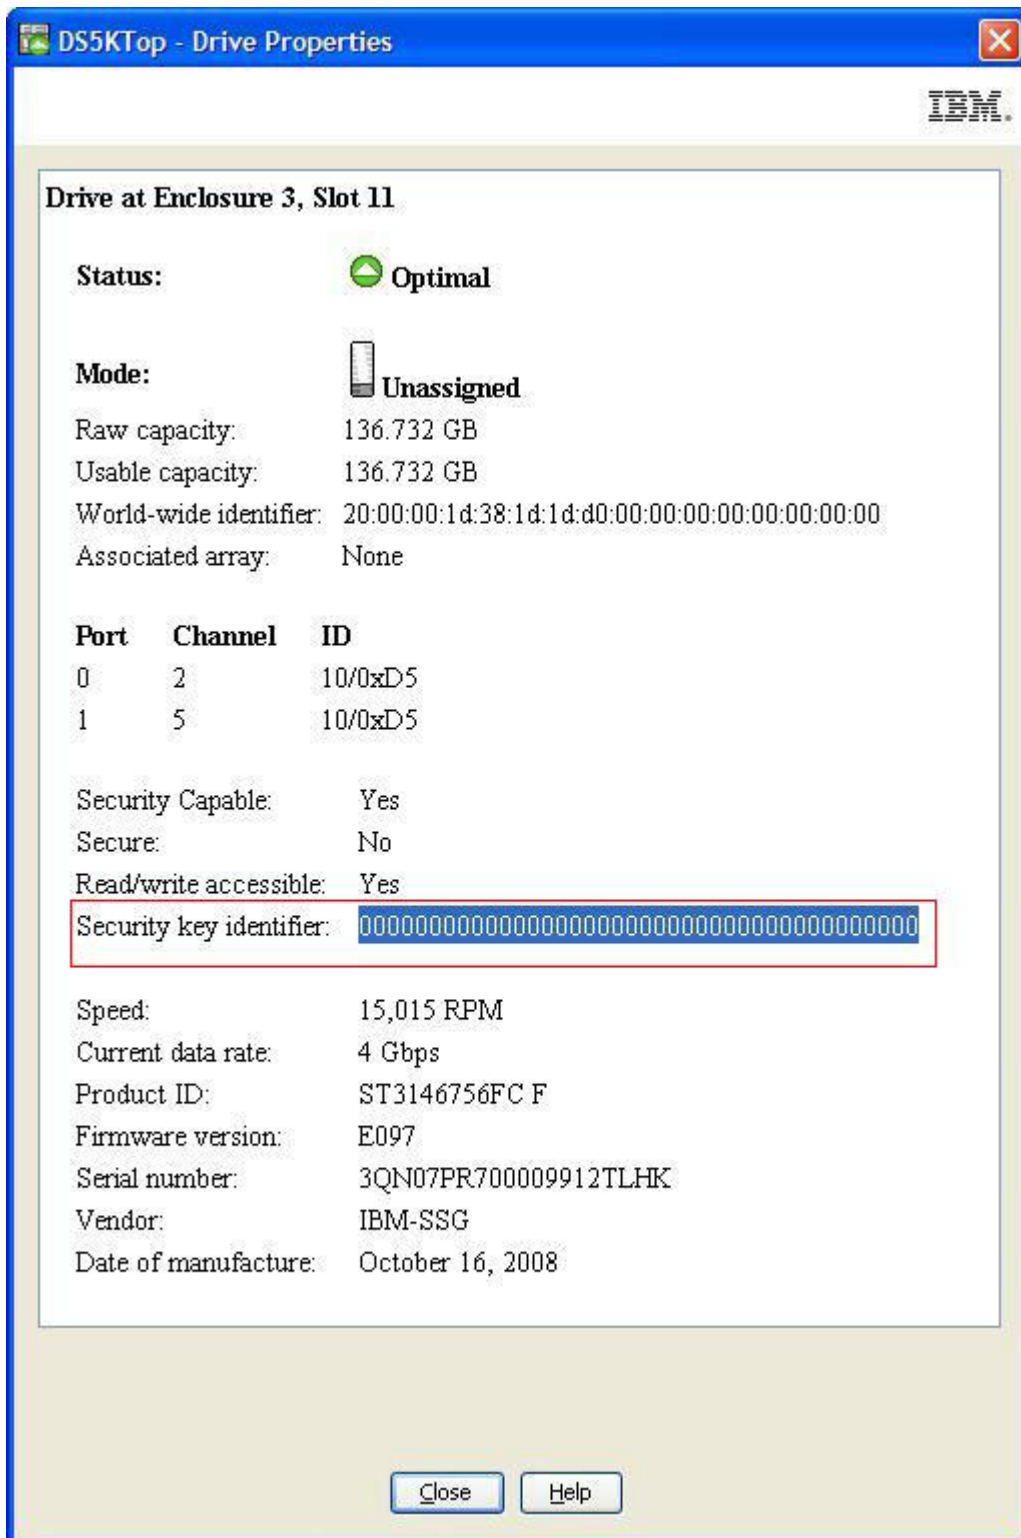


Figure 31. Drive properties - Unsecured FDE drive

Unlocking secure drives in local security key management mode

You can export a RAID array with its security-enabled FDE drives to a different storage subsystem. After you install those drives in the new storage subsystem,

you must unlock the security-enabled FDE drives before data can be read from or written to the drives. The security key on the new storage subsystem will be different and will not unlock the drives. You must supply the security key from a security key file that you saved from the original storage subsystem. In addition, you must provide the pass phrase that was used to encrypt the security key to extract the security key from the security key file. After you unlock the drives with the security key in the security key file, the controller negotiates the existing security key for these drives so that only one version of the security key is used to unlock drives in a storage subsystem.

You do not have to provide the security key file to unlock the security-enabled drives in a storage subsystem every time the storage subsystem power is cycled or the drives are removed and reinserted in the same storage subsystem, because the controllers always keep a copy of the current and previous ($n-1$) values of the security key to unlock these drives. However, if the drives are removed from the storage subsystem and the security key is changed more than two times in the same storage subsystem, the controllers will not have the security key to unlock the drives when they are reinserted in the same storage subsystem.

Attention: Always back up the data in the storage subsystem to secured tape to prevent loss of data due to malicious acts, natural disasters, abnormal hardware failures, or loss of the FDE security key.

Unlocking secure drives in external security key management mode

You can export a RAID array with its security-enabled FDE drives to a different storage subsystem. If the drives are moving to another subsystem that is managed by the same external key license manager as the original subsystem, you do not have to provide the saved security key file to unlock the drives.

After you unlock the drives with the security key in the security key file, the controller negotiates the existing security key for these drives so that only one version of the security key is used to unlock drives in a storage subsystem. Otherwise, you must supply the security key from a security key file that you saved from the original storage subsystem. You must also provide the pass phrase that was used to encrypt the security key to extract the security key from the security key file. After you unlock the drives with the security key in the security key file, the controller negotiates the existing security key for these drives so that only one version of the security key is used to unlock drives in a storage subsystem.

Note: You must export the array from the original subsystem before you move the FDE drives; the array is required to configure the drives for removal and update the subsystem configuration.

If the subsystem configuration does not have any unsecured drives or non-FDE drives, you must provide the security key from a security key file that you saved from the original storage subsystem when the power to the new subsystem is turned on. If the subsystem configuration has optimal unsecured or non-FDE drives, the subsystem will start up and connect with the external security key manager to obtain the key that unlocks the secured FDE drives.

Attention: Always back up the data in the storage subsystem to secured tape to prevent loss of data due to malicious acts, natural disasters, abnormal hardware failures, or loss of the FDE security key.

Using secure erase

Secure erase protects FDE drives from security threats when they are eventually retired, returned, discarded, or re-purposed. As these drives are moved from the data center or reused, it is critical that the data on the disks be permanently erased and not vulnerable to recovery. Discarded drives might still have residual data that can be reconstructed by an unauthorized user. Secure erase protects against this threat by cryptographically erasing the data.

The traditional methods that are used to permanently erase data often prove to be expensive and slow and might not provide the highest level of data erasure. Traditional methods might also put the drives beyond your control and therefore subject to a data breach. Secure erase provides the following advantages compared to traditional methods:

- Immediate, cryptographic data erasure
- Lower overall costs
- A higher level of media sanitation, in accordance with the National Institute of Standard and Technology (NIST)

Attention: Secure-erase operations are not reversible. All data on the drive will be permanently erased when a secure-erase action is performed. Make sure that the data on the drive is backed up or expendable.

Secure erase with FDE drives allows for immediate erasure of data without requiring that the drive be removed from the data center. With just a few clicks, you can quickly reuse or discard a drive. With secure erase, you can erase drives and use them again. This eliminates the need to destroy a drive, yet still secures warranty and expired lease returns and enables you to reuse drives securely. According to the NIST, secure erase is considered a type of data purging, which is regarded as a higher level of data sanitation than traditional methods.

Secure erase prompts the FDE drive to permanently erase the current encryption key and replace it with a new randomly-generated encryption key within the drive. The drive encryption key is used to encode and decode all data on the disk. After the encryption key is changed, any data that was previously written to the disk becomes unintelligible. Data that was encrypted with the previous encryption key is unintelligible when it is decrypted with the new encryption key. This includes all bits, headers, and directories. The data is completely and permanently inaccessible.

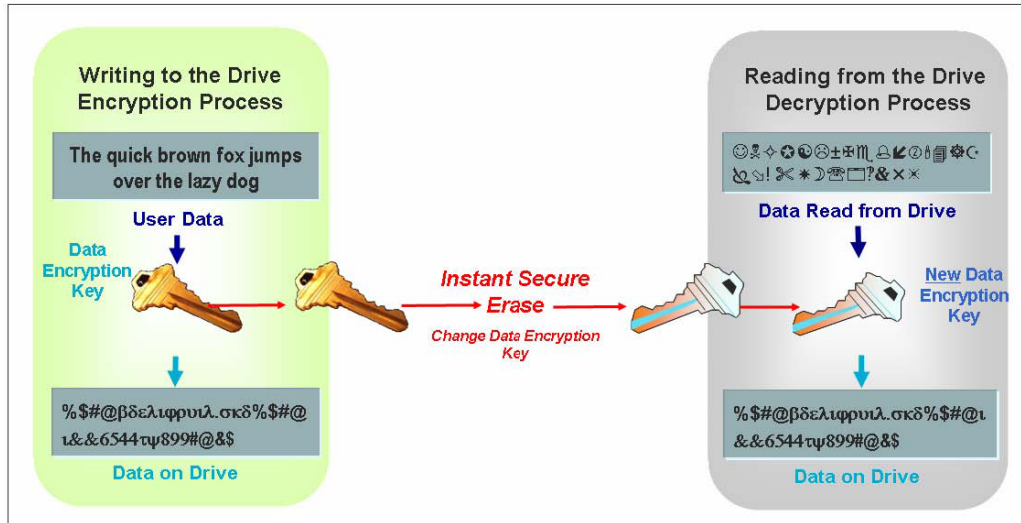


Figure 32. Secure erase process

FDE security authorizations

The following table identifies and describes the authorization parameters that are used to implement security on FDE-compatible storage subsystems.

Table 35. Security authorizations

Parameter	Description	Where is it located and managed?	How is it generated?
Encryption key	The encryption key is used to encrypt and decrypt data on the FDE disk drive.	Is stored on and is managed by the FDE disk drive: <ul style="list-style-type: none"> It is never transferred from the drive. Each drive has its own unique encryption key. 	The encryption key is generated when the drive is manufactured and then regenerated at the customer site (by a command from the controller to the drive) to ensure that the key was not compromised prior to use.
Security key	The security key is needed to unlock the encryption key for encrypting and decrypting to occur. One security key is created for all FDE drives on the storage subsystem. The security is sometimes referred to as the lock key.	Is stored on and is managed by the controller. A single security key is synchronized for all controllers in a storage subsystem.	The security key is generated by the storage subsystem and is encrypted and hidden in the storage subsystem.

Table 35. Security authorizations (continued)

Parameter	Description	Where is it located and managed?	How is it generated?
Security key identifier	The security key identifier is paired with the security key to help you remember which key to use for secure operations. With local security key management only, you have the option to provide up to 189 alphanumeric characters that are linked to the storage subsystem-generated security key identifier.	The security key identifier is stored in a special area of the disk: <ul style="list-style-type: none"> • Can always be read from the disk • Can be written to the disk only if security has been enabled and the drive is unlocked 	User-specified alphanumeric character string (local security key management only). The storage subsystem adds the storage subsystem worldwide identifier and a randomly generated number to the characters that are entered.
Pass phrase	The pass phrase is used to encrypt the security key and the security key identifier. The pass phrase is a user-specified alphanumeric character string, eight characters minimum, 32 characters maximum. It must contain at least one number, one lowercase letter, one uppercase letter, and one nonalphanumeric character (such as <, >, &, @, +, or -). Spaces are not allowed, and it is case-sensitive.	User-specified alphanumeric character string, not stored anywhere on the storage subsystem or in the security key file. The pass phrase is used to encrypt the security key when it is exported in the security key file. It is also used to decrypt the key in the security file when it is used to import security-enable FDE drives into a storage subsystem.	User-specified alphanumeric character string.
Security key file	File where the security key identifier is saved along with the encrypted security key.	File name and location are determined by the administrator. In addition to the administrator-specified location, the storage manager also saves a copy of the security key backup file in the default location. See the <i>IBM Full Disk Encryption Best Practices</i> document for more information.	Generated by the storage subsystem after you initiate a create security key, change security key, or save security key operation.

FDE terminology

The following table defines FDE terminology that is used throughout this chapter.

Table 36. Full disk encryption terminology

Term	Description
FDE	Full disk encryption, a custom chip or ASIC (application specific integrated circuit) on the disk drive that requires a security key to allow encryption and decryption to begin. FDE disk drives encrypt all the data on the disk. The secured drive requires that a security key be supplied before read or write operations can occur. The encryption and decryption of data is processed entirely by the drive and are not apparent to the storage subsystem.
Secure erase	Permanent destruction of data by changing the drive encryption key. After secure erase, data that was previously written to the disk becomes unintelligible. This feature takes advantage of FDE disk security capabilities to erase data by changing the encryption key to a randomly generated value. Because the encryption key never leaves the drive, this provides a secure erase. After secure erase, the drive becomes unlocked, allowing anyone to read or write to the disk. Secure erase is sometimes referred to as drive reprovisioning.
Local key management	A key management method that uses a security key created and contained in the storage subsystem controller. To move secured drives from one storage subsystem to another, you must use the saved security key file from the original storage subsystem to unlock the drives. The security key is obfuscated and stored in the storage subsystem when the power is turned off.
External key management	A key management method that uses a central key location on your network (one or more servers external to a storage subsystem) to manage keys for different storage devices. A proxy server must facilitate the request for and acceptance of a security key. The security key is not stored in the storage subsystem when the power is turned off. Note: 1. External security key management requires dedicated software, such as IBM Tivoli Key Lifecycle Manager (TKLM).
Locked	The state that a security-enabled FDE drive enters when it has been removed from and then reinserted in the storage subsystem, or when the storage subsystem powered off. When storage subsystem power is restored, the drive remains in the Locked state. Data cannot be written to or read from a locked disk until it is unlocked by the controller, using the security key. If the controller does not have the security key, the security key file and its pass phrase are required to unlock the drives for read and write operations.
Repurposing/Reprovisioning	Changing a drive from being in Secured state to Unsecured state so that the drive can be reused. Reprovisioning the drive is accomplished by secure erase.
Secure array	An array on security-enabled FDE drives.
Security-capable drive	An FDE drive that is capable of encryption but is in Unsecured state (security not enabled).
Security-enabled drive	An FDE drive with security enabled. The security-enabled FDE drive must be unlocked using the security key after power to the drive is turned on and before read or write operations can occur.
Unlocked	The state of a security-enabled FDE drive in which data on the disk is accessible for read and write operations.

Before you begin

If you use external security key management, you must complete the following procedures:

1. Install and configure the external key license manager software, IBM Tivoli Key Lifecycle Manager (TKLM). See the documentation that came with the software for more information.
2. Download the DS TKLM Proxy Code from the IBM Support Portal at <http://www.ibm.com/support/entry/portal>.
3. Install and configure the DS TKLM Proxy Code. See “Installing and configuring the DS TKLM Proxy Code server.”
4. Enable the Full Disk Encryption and External Key Management premium features in Storage Manager. See “Enabling premium features” on page 217.
5. Configure TKLM and the storage subsystems for the DS TKLM proxy and create external key management security authorizations. See “Creating security authorizations using external security key management” on page 220 in “Enabling premium features” on page 217.

If you prefer to use local security key management, begin with the information in “Configuring disk encryption with FDE drives” on page 216.

Installing and configuring the DS TKLM Proxy Code server

This section describes the procedures that are required to install the DS TKLM Proxy Code server. The DS TKLM Proxy Code supports the following operating systems:

- AIX 5.x
- AIX 6.x
- Red Hat Enterprise Linux 4.x
- Red Hat Enterprise Linux 5.5
- SUSE Linux Enterprise Server 10.3
- SUSE Linux Enterprise Server 11
- Windows 2008 R2
- Windows 2008 Service Pack 2
- Windows 2008 Standard Edition
- Windows 2008 Enterprise Edition

Important: Any environmental or configuration change that might affect the DS TKLM Proxy Code server requires that you restart the server. For example, a Storage Manager controller swap, the issuing of the **sysWipe** command, or a change to the IP address would require that the DS TKLM Proxy Code server be reconfigured and restarted. In addition, a change to the security key identifier, such as swapping the Storage Manager controller or issuing the **sysWipe** command, requires that TKLM be modified to recognize the new security key identifier as defined in “Creating security authorizations using external security key management” on page 220. See “Starting, stopping, and restarting the DS TKLM Proxy Code server” on page 211 for more information.

The following diagram illustrates the relationships between the components of an external security key management configuration.

Note:

1. A maximum of four storage subsystem controllers can be monitored by one proxy server.
2. A maximum of four TKLM servers can be connected to one proxy server.

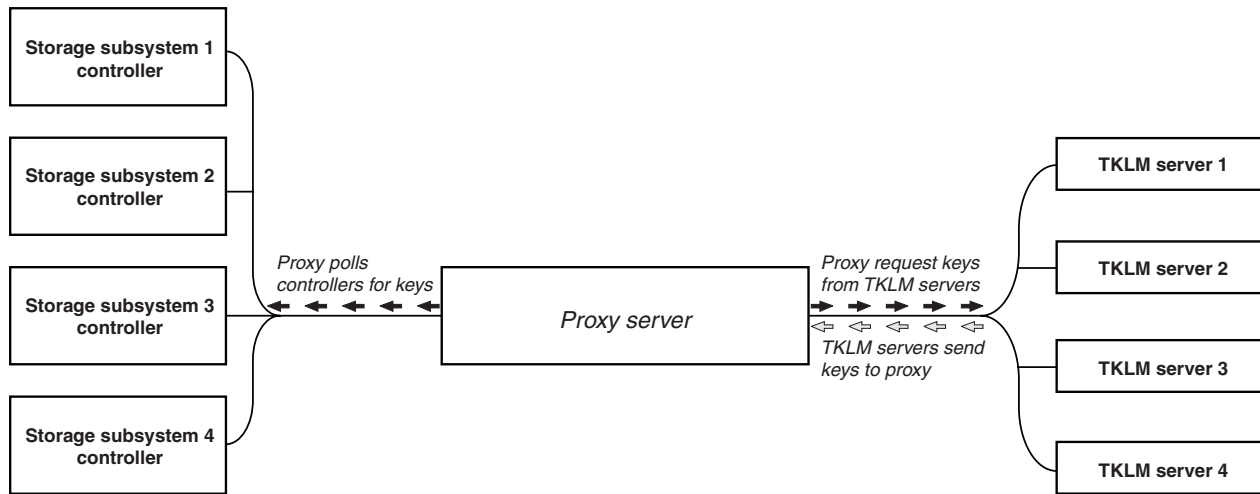


Figure 33. External security key management topology

To establish an external security key management configuration, download the DS TKLM Proxy Code from the IBM Support Portal at <http://www.ibm.com/support/entry/portal> and complete the following procedures:

1. "Modifying the DS TKLM Proxy Code server configuration file" on page 212
2. "Installing the DS TKLM Proxy Code" on page 215

Important: You must complete the procedures in order. Make sure that the IBM Tivoli Key Lifecycle Manager (TKLM) software is installed. See the documentation that came with the software for more information.

Starting, stopping, and restarting the DS TKLM Proxy Code server

Any environmental or configuration changes that might affect the proxy (for example, network changes) require the proxy to be restarted. The following utilities are available.

For AIX:

```
start_DS_TKLM_Proxy_Code_AIX.sh
stop_DS_TKLM_Proxy_Code_AIX.sh
restart_DS_TKLM_Proxy_Code_AIX.sh
```

For Linux:

```
start_DS_TKLM_Proxy_Code_Linux.sh
stop_DS_TKLM_Proxy_Code_Linux.sh
restart_DS_TKLM_Proxy_Code_Linux.sh
```

The `stop_DS_TKLM_Proxy_Code_*.sh` script will remove the entry from `/etc/inittab` and end the processes.

Modifying the DS TKLM Proxy Code server configuration file

The configuration file for the proxy is `DS_TKLM_Proxy_Code.config`. The configuration file name, as well as the parameters and their definitions, are the same for all supported operating systems (Windows, AIX, and Linux). However, the format of some of the parameter values are different in Windows and AIX or Linux.

The method of creating and editing the configuration file in Windows is different from the method in AIX or Linux. With Windows, you must create `DS_TKLM_Proxy_Code.config` manually, using the template included in the `DS_TKLM_Proxy_Code_Windows*.zip` file. The definitions for parameters must be assigned before the proxy can be installed.

Important: If you are working in a Windows operating-system environment, you must create and modify the configuration file before you install the DS TKLM Proxy Code server.

With AIX and Linux, `DS_TKLM_Proxy_Code.config` is created and parameter definitions are assigned during the installation. You must assign definitions for the configuration file parameters when you are prompted.

The definition of each parameter is explained in the following table.

Table 37. Proxy configuration file properties

Property name	Description	Example
LogLevel	This property specifies one of the following four levels for logging: <ul style="list-style-type: none">• UserInfo: basic information about events• UserWarning: warning information about a potential problem• UserError: error information about a system failure, the proxy server exits• Debug: information useful for debugging, such as string lengths and property values on different positions	LogLevel = debug
DebugPath	This property specifies the location of the debug file. You must provide a path in your file system that can either be a path relative to the directory <code>/DS_TKLM_Proxy_Code/bin</code> or an absolute path. Note: Make sure that you have read and write permissions for the path directory.	AIX or Linux example: DebugPath = <code>./Log/Debug/debug.log</code> Windows example: DebugPath = <code>.\Log\Debug\debug.log</code>
AuditPath	This property specifies the location of the audit file. You must provide a path in your file system that can either be a path relative to the directory <code>/DS_TKLM_Proxy_Code/bin</code> or an absolute path. Note: Make sure that you have read and write permissions for the path directory.	AIX or Linux example: AuditPath = <code>./Log/Audit/audit.log</code> Windows example: AuditPath = <code>.\Log\Audit\audit.log</code>

Table 37. Proxy configuration file properties (continued)

Property name	Description	Example
ThresholdSize	This property specifies the maximum size of each log file in bytes. If the size threshold is reached, a new file is created with same file name as the original file name with the numbers 01 added at the end. If the new log file reaches the size threshold, the original file is overwritten. Note: If you decide later to increase the threshold size, delete the existing log files. Otherwise, the proxy will write log information in the old files if the new size threshold is larger than the old size threshold.	Threshold size = 100000000000
KeyinformationPath	This property specifies the location of the security certificate file (matched with the file specified in KeyPassword property). Enter a path in your file system that can be either relative to /DS_TKLM_Proxy_Code/bin or an absolute path. Make sure that the path adheres to conventions for directory specification for Windows or AIX and Linux, and make sure that the directory and file that you specify exist. Note: This property refers to the security certificate file and password you received in an email after you enabled the External Key Management premium feature. If you do not receive the security certificate file or if you no longer have the file, you can request another file and password by using the key reactivation process on the IBM Premium Features website.	AIX or Linux example: KeyinformationPath = ./CertFile/ibmproxycert.p12 Windows example: KeyinformationPath = .\CertFile\ibmproxycert.p12
KeyPassword	This property specifies the password for the security certificate (matched with file specified in the KeyinformationPath property), and will be obfuscated after reading occurs. If the password must be changed after it has been obfuscated, you must first delete the KeyPasswordHex property value and restart the proxy server. Otherwise, the new password is ignored. Note: This property refers to the security certificate file and password you received in an email after you enabled the External Key Management premium feature. If you do not receive the security certificate file or if you no longer have the file, you can request another file and password by using the key reactivation process on the IBM Premium Features website.	Example of KeyPassword property before the first reading occurs: KeyPassword = password Example of KeyPassword property after first reading occurs and the password is obfuscated: KeyPasswordHex = 47558BADDI3321FC KeyPassword = *****

Table 37. Proxy configuration file properties (continued)

Property name	Description	Example
SYMServer.x	<p>The term <i>SYMServer</i> refers to a storage subsystem and its controllers.</p> <p>Note: A maximum of four storage subsystem controllers can be monitored by one proxy server. This property specifies information about every storage subsystem, or Symbol-Server (SYMServer.1 - SYMServer.n) in your configuration. Each SYMServer requires two controller IP addresses, two ports (2463), one SSID, one password indicator, and one password. Therefore, each SYMServer property value must match the pattern in the following format. The variables are italicized:</p> <p><i>SYMSEVER.x = Controller A IP address , Controller B IP address , port number , port number , SSID , password indicator , password</i></p> <p>The password indicator must be set to false if the password is provided in clear text and true if the password is obfuscated. The password is used to manage the storage subsystem. The password will be obfuscated automatically and stored in an encrypted format by the proxy after reading occurs.</p> <p>In this property, you can use spaces between each part of the value. The SSID has to be a hexadecimal value. The proxy compares the SSID from the DS_TKLM_Proxy_Code.config file to the SSID that it retrieves from the storage subsystem. If they are not equal, the proxy will stop monitoring the storage subsystem.</p> <p>Note: You must obtain the SSID for the storage subsystem from the Storage Manager Storage Subsystem Profile window.</p>	<p>Example before the first time the configuration file is read:</p> <p>SYMServer.1 = 9.37.117.35 , 9.37.117.36 , 2463 , 2463 ,600A0B8000339848000000004B72851F, false, SymPasswd</p> <p>Example after the first time the configuration file is read:</p> <p>SYMServer.1 = 9.37.117.35 , 9.37.117.36 , 2463 , 2463 , 600A0B8000339848000000004B72851F, true , 6408D5D0C596979894AA8F</p>
TKLMServer.x	<p>This property specifies information about every TKLM server in your configuration.</p> <p>Note: A maximum of four TKLM servers can be connected to one proxy server. Each TKLM server has one IP address and one port, so each TKLM server property value must match the pattern in the following format. The variables are italicized:</p> <p><i>TKLMServer.x = IP address , port number</i></p> <p>In this property, you can use spaces between each part of the value. If you do not enter a value for this property, the proxy server uses the default value (localhost, 3801). The port number is found on the Key Serving Ports window in the Tivoli Lifecycle Key Manager software.</p>	TKLMServer.1 = 9.41.18.161 , 3801
TcpTimeout	This property specifies the length of the timeout period for a TCP connection to the servers, in seconds.	TcpTimeout = 1000
RpcTimeout	This property specifies the length of the timeout period for remote procedure calls on servers, in seconds.	RpcTimeout = 10
TimeBetween-SymbolServer-Queries	This property specifies a waiting period between proxy-server checks of the attention state, in seconds.	TimeBetweenSymbolServerQueries = 10

Installing the DS TKLM Proxy Code

To install the DS TKLM Proxy Code for use with external security key management, complete one of the following procedures. For a Windows environment, see “Installing the DS TKLM Proxy Code server in a Windows environment.” For an AIX or Linux environment, see “Installing the DS TKLM Proxy Code server in an AIX or Linux environment.”

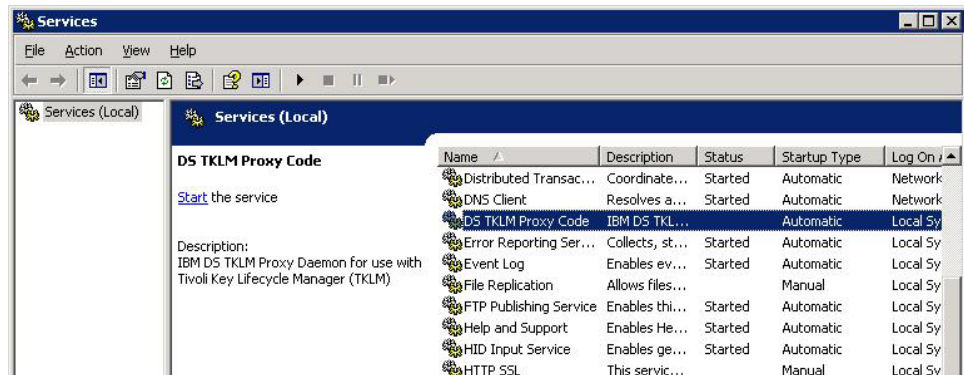
Installing the DS TKLM Proxy Code server in a Windows environment

To install the proxy for a Windows environment, complete the following steps:

1. Go to the IBM Support Portal at <http://www.ibm.com/support/entry/portal> and download the applicable file for your version of Windows from the Downloads and fixes area of the portal. The file name is similar to DS_TKLM_Proxy_Code-Windows-V*.zip.
2. Extract the compressed files to a local directory (for example, c:\DS_TKLM_Proxy_Code).
3. Make sure that the DS_TKLM_Proxy_Code.config file has been modified (see “Modifying the DS TKLM Proxy Code server configuration file” on page 212 for the properties that must be modified).
4. Make sure that the certificate file, obtained from IBM and specified in the KeyInformationPath property in the configuration file, exists before you start the proxy server.

Note: If a “DS_TKLM_Proxy_Code_WinService.exe - Application Error” message is displayed, you might have to download and install the Microsoft Visual C++ Redistributable Package. For the package that is compatible with Windows 2008, go to <http://www.microsoft.com/downloads/details.aspx?familyid=A5C84275-3B97-4AB7-A40D-3802B2AF5FC2&displaylang=en>.

5. In a DOS prompt window, type the following command:
DS_TKLM_Proxy_Code_WinService.exe -i. The proxy is added to the Services window. To start the proxy, click **Start** in the Services window.



Note: To uninstall the proxy, open a DOS prompt window and type and execute the following command: DS_TKLM_Proxy_Code_WinService.exe -u. Restart Windows.

Installing the DS TKLM Proxy Code server in an AIX or Linux environment

The DS TKLM Proxy Code is packaged in RPM format for AIX or Linux (RedHat and SUSE). To install the proxy server in an AIX or Linux environment, complete the following steps:

1. Go to the IBM Support Portal at <http://www.ibm.com/support/entry/portal> and download the applicable file for your operating system version from the Downloads and fixes area of the portal. For example, the file name for AIX might be DS_TKLM_Proxy_Code-AIX-V2.01_90.70.G0.04.ppc.rpm, and the file name for Linux might be DS_TKLM_Proxy_Code-Linux-V2.01_90.70.G0.04.i386.rpm.

Note: Be sure to download the correct file for your operating system. The operating system is a part of the RPM file name.

2. Use rpm commands to extract the downloaded file and begin the installation process. For example:

```
rpm -ivh --nodeps DS_TKLM_Proxy_Code-AIX-V1_.ppc.rpm
```

Note: The --nodeps part of the command is required only for AIX installations. When you execute the RPM command, you create symbolic links, specify the location of the certificate file that is provided by IBM, create a backup of /etc/inittab, and provide the path to use when you execute the installation script.

3. After you execute the RPM command, run the installation script (/DS_TKLM_Proxy_Code/bin/install.sh).
4. When you are prompted, enter all of the configuration file properties. See “Modifying the DS TKLM Proxy Code server configuration file” on page 212 for a description of the properties and their values.

To configure TKLM and storage subsystems for the proxy, and to create external key management security authorizations, continue to “Creating security authorizations using external security key management” on page 220 in “Enabling premium features” on page 217.

Configuring disk encryption with FDE drives

This section provides the procedures for enabling FDE and creating secure arrays on the storage subsystem. To configure disk encryption with FDE disks, perform the following tasks:

1. Install the FDE drives (see “Installing FDE drives” on page 217).
2. Enable the Full Disk Encryption premium feature (see “Enabling premium features” on page 217).
3. Create an array and enable array security (see “Securing a RAID array” on page 224). You can also enable security for a disk pool. The procedure is exactly the same as that for an array.

Note: The screenshots in this section are only illustrative and may differ from the actual UI depending on the Storage Manager and controller firmware versions.

A security-enabled FDE drive becomes locked when its power is turned off or when it is removed from the storage subsystem. To unlock a locked drive, see “Unlocking disk drives” on page 229.

In some storage subsystems, drives can be migrated as a complete array into another storage subsystem. To migrate a secure array, see “Migrating storage subsystems (head-swap) with FDE drives” on page 232.

Installing FDE drives

This section lists the FDE disk drives that FDE-compatible IBM DS storage subsystems support, as of the date of this document. See the *IBM System Storage DS3000, DS4000, or DS5000 Hard Drive and Storage Expansion Enclosure Installation and Migration Guide* and the *DS5000 Interoperability Guide* for installation procedures and the most up-to-date support information.

The FDE drives that are supported by an IBM DS storage subsystem are specified in the announcement letters for that particular storage subsystem. You can also contact your IBM reseller or IBM marketing representative for more information about compatible FDE drives for your storage subsystem.

Note: If the FDE drive is in Security Enabled state and you do not want to preserve the data on the drive, perform a secure erase on each drive before you use it as part of a new RAID array. Secure erase forces the drive to generate a new encryption key, places the drive in Unsecured state, and ensures that any data that was previously stored on the disk is erased. See “Using secure erase” on page 206 for more information.

Enabling premium features

The FDE premium feature must be enabled on the storage subsystem, using the instructions that come with the IBM DS Disk Encryption premium feature key entitlement kit. To verify that full disk encryption is enabled, on the Setup page, select **View/Enable Premium Features**. In the Premium Features and Feature Pack Information window, Full Disk Encryption: Enabled and External Key Management: Enabled indicates that the FDE premium feature is enabled.

Important: External key management requires a security certificate file and its password. The file and password are emailed to you after you enable the External Key Management premium feature. When you enable the External Key Management premium feature at the IBM Premium Feature website, you must provide a valid email address in the fields shown in the following image. Otherwise, you are prompted to enter your email address after you click **Continue**.

Your activation key file will be provided via a link once the information above is verified and submitted. In addition, if you would like the activation key file sent to you, please provide your email address below:

Email address

Verify email address

E-mail: Stay informed about IBM products, services, and other offerings! If you want to stay informed by e-mail, please let us know by checking the box below.

☐ e-mail: Yes, please have IBM or an affiliate send me e-mail.

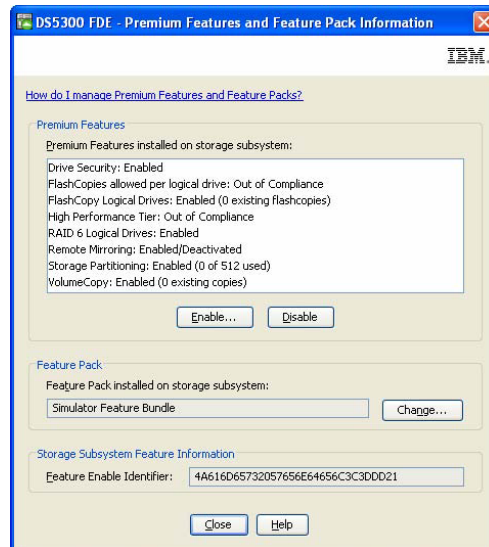
Other communications: IBM or an affiliate or selected organizations may keep you informed about IBM related products, services and other offerings through ways other than e-mail, for example, by telephone or postal mail. If you do not want us to use the information you provided here to keep you informed through other ways, please indicate in the box below.

☐ Other communications: Please do not use the information I have provided here.

By clicking "Continue", you agree that IBM may process your data in the manner indicated above and as described in our Privacy policy.

It might take up to a day to receive the security certificate file and password. If you do not receive the file or if you no longer have the email with the file, you can request another file and password by using the key reactivation process on the IBM Premium Features website. For more information about the security certificate file and configuring the KeyInformationPath and KeyPassword properties (Windows operating systems only), see “Modifying the DS TKLM Proxy Code server configuration file” on page 212.

If you enable the FDE feature after November, 2010, for a storage subsystem with controller firmware 7.70.xx.xx or later, External Key Management: Enabled and Full Disk Encryption: Enabled are displayed in the Premium Features and Feature Pack Information window.



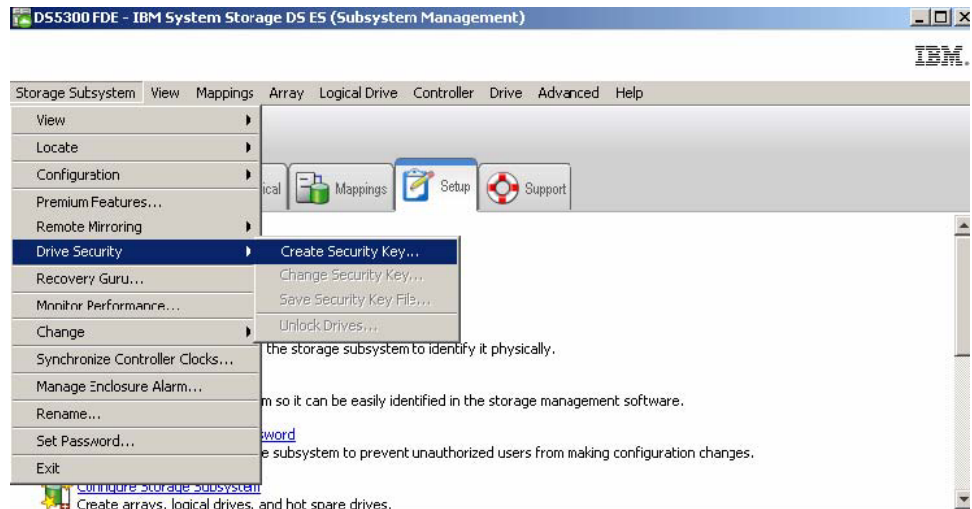
Enabling full disk encryption includes creating the security authorizations that you will need later to unlock a secured FDE drive that has been turned off or removed from the storage subsystem. These authorizations include the security key identifier, a pass phrase, and the security key file. The security authorizations apply to all the FDE drives within the storage subsystem and are critical if a drive must be unlocked after the power is turned on.

The process for creating security authorizations depends on the method of key management you use. See the applicable section for local or external security key management.

Creating security authorizations using local security key management

To create the security authorizations for full disk encryption using local key management, complete the following steps. For external key management security authorizations, see “Creating security authorizations using external security key management” on page 220.

1. From the Storage Manager Subsystem Management window, click **Storage Subsystem**, click **Drive Security**, and click **Create Security Key**.



2. Enter a security key identifier, the security key file name and location, and a pass phrase in the Create Security Key window:
 - **Security key identifier:** The security key identifier is paired with the storage subsystem worldwide identifier and a randomly generated number and is used to uniquely identify the security key file. The security key identifier can be left blank or can be up to 189 characters.
 - **Pass phrase:** The pass phrase is used to decrypt the security key when it is read from the security key file. Enter and record the pass phrase at this time. Confirm the pass phrase.
 - **Security key backup file:** Click **Browse** next to the file name to select the security key file name and location, or enter the value directly in the field. Click **Create Key**.

Note: Save the security key file to a safe location. The best practice is to store the security key file with your key management policies. It is important to record and remember where this file is stored because the security key file

- is required when a drive is moved from one storage subsystem to another or when both controllers in a storage subsystem are replaced at the same time.
3. In the Create Security Key Complete window, record the security key identifier and the security key file name; then, click **OK**. The authorizations that are required to enable security on FDE drive in the storage subsystem are now in place. These authorizations are synchronized between both controllers in the storage subsystem. With these authorizations in place, arrays on the FDE drives in the storage subsystem can be secured.

Attention: For greater security, store more than one copy of the pass phrase and security key file. Do not specify the default security file directory as the location to store your copy of the security key file. If you specify the default directory as the location to save the security key file, only one copy of the security key file will be saved. Do not store the security key file in a logical drive that is mapped from the same storage subsystem. See the *IBM Full Disk Encryption Best Practices* document for more information.




Creating security authorizations using external security key management

Important:

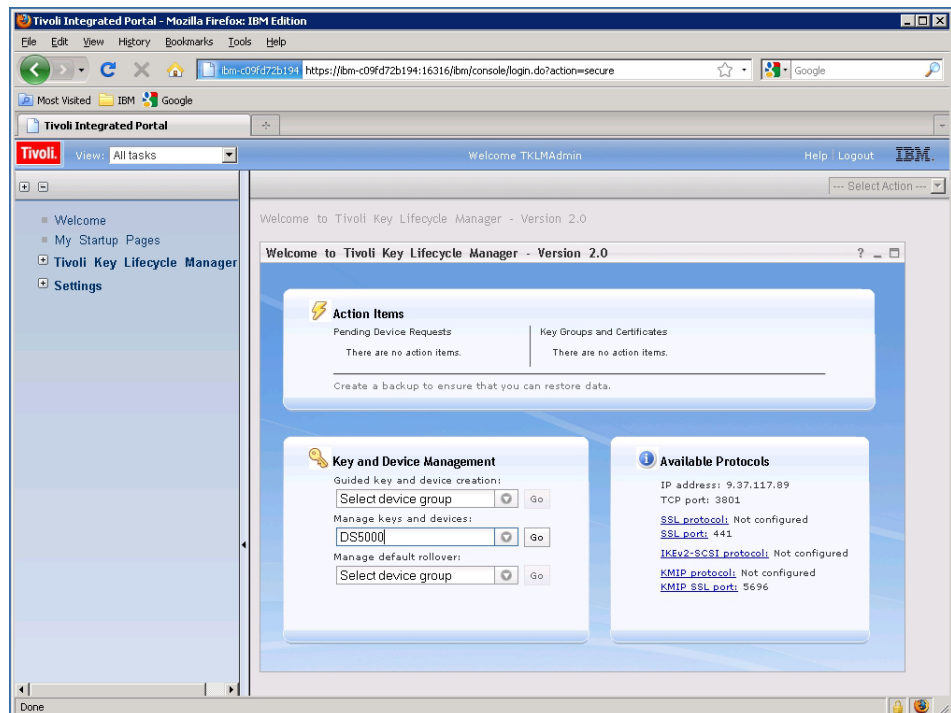
Before you can create the security authorizations for full disk encryption with external key management, you must complete the procedures in “Installing and configuring the DS TKLM Proxy Code server” on page 210.

After the Tivoli Key Lifecycle Manager (TKLM) is installed, it must be configured to service key requests from the DS TKLM Proxy Code server. To configure TKLM, complete the following steps:

1. Open TKLM and log in with the TKLAdmin ID.
2. Click **click here to create the master keystore**. The Keystore settings window is displayed.
3. Type and retype the password for the keystore. Keep the default values for the other keystore settings, and click **OK**.



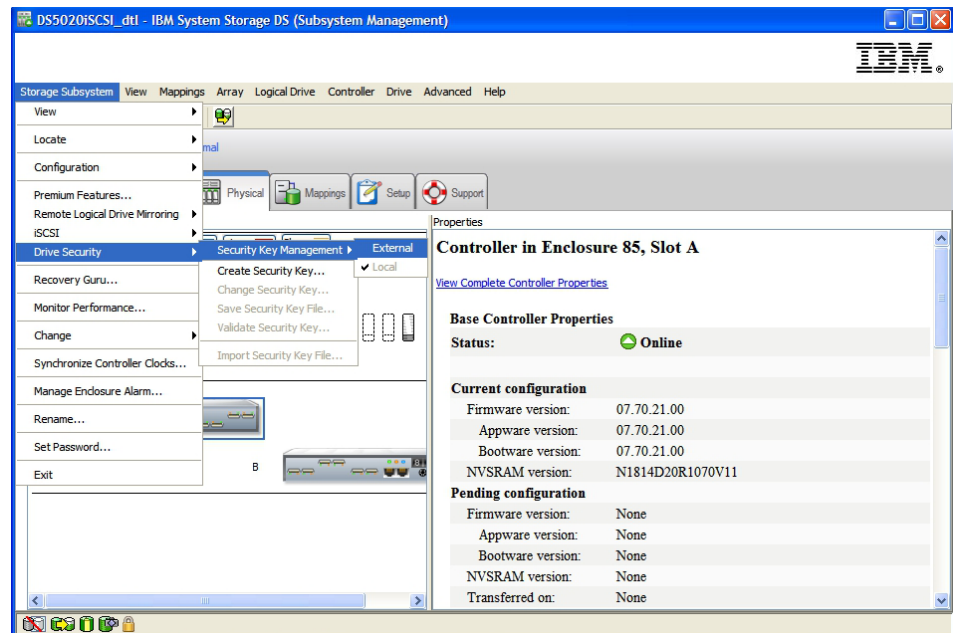
4. Click the **Welcome** link on the left side of the window. The Welcome window opens.
5. In the **Key and Device Management** box, select **DS5000** from the **Manage keys and devices** menu, and click **Go**. The Key and Device Management window opens.



6. When the Confirm prompt is displayed, click **Cancel**.

7. In the drop-down menu at the bottom of the window, select **Hold new device requests pending my approval**.
8. Open Storage Manager, log in, and open the Subsystem Management window for the storage subsystem that you are configuring.
9. Click **Storage Subsystem > Drive Security > Security Key Management > External**.


Note: If the External Key Management premium feature is not enabled, the menu option **Security Key Management** is not displayed when you click **Storage Subsystem > Drive Security**.




10. When you are prompted with the Confirm Security Key Management window, type yes and click **OK**.



11. When you are prompted, save a copy of the security key. Enter the pass phrase, file name, and file location, and click **OK**. The controller attempts to contact the external key manager for the security key. If it fails, the following message is displayed:

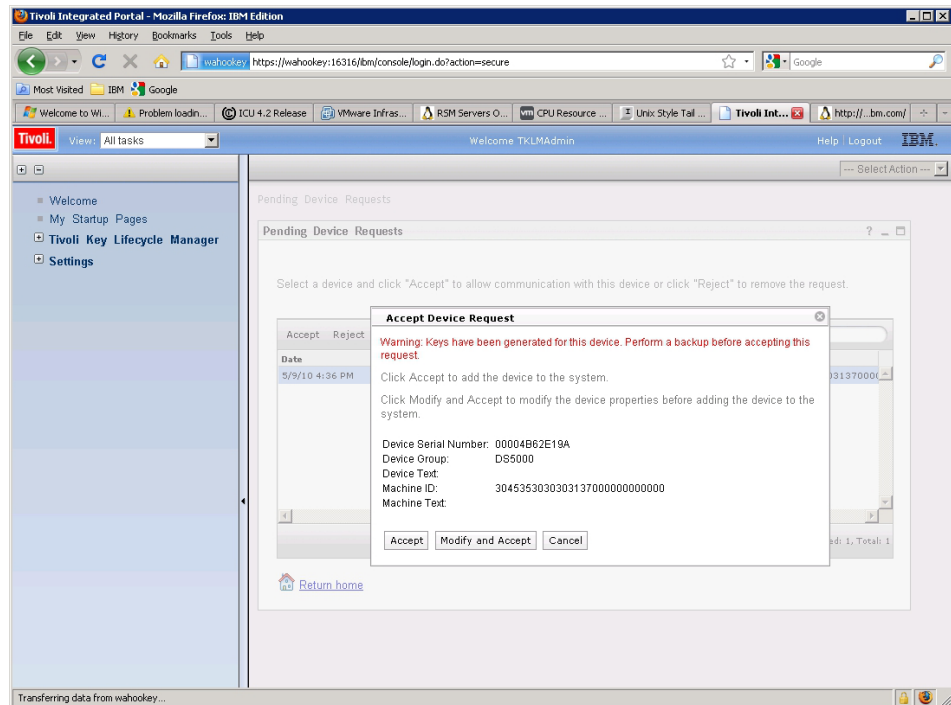


12. Return to the TKLM application and click the **Pending devices** link in the **Action Items** box.



The Pending Device Request window opens.

13. Select the device in the list and click **Accept**. The Accept Device Request window opens.
14. Click **Accept** on the Accept Device Request window.



The TKLM server is now ready to send keys to the DS TKLM Proxy Code server.

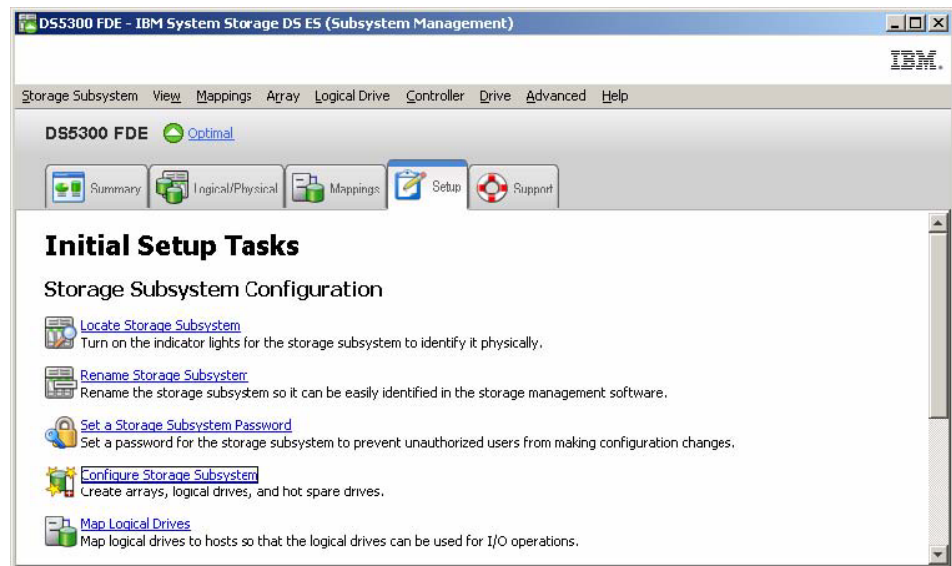
Securing a RAID array

An array is secured when the FDE drives in the array are security enabled. The FDE drives in a secured array become locked if their power is turned off or if they are removed from the storage subsystem.

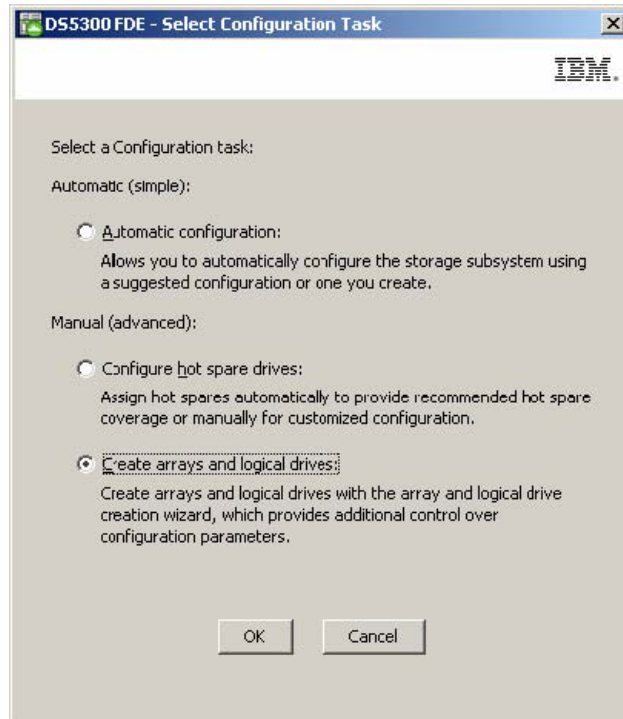
All drives in the array must be security-capable FDE drives with security not enabled. The array must not contain any FlashCopy base logical disks or FlashCopy repository logical disks. Base logical disks and FlashCopy logical disks can be written to the disks only after security is enabled.

To create a RAID array and then secure it, complete the following steps:

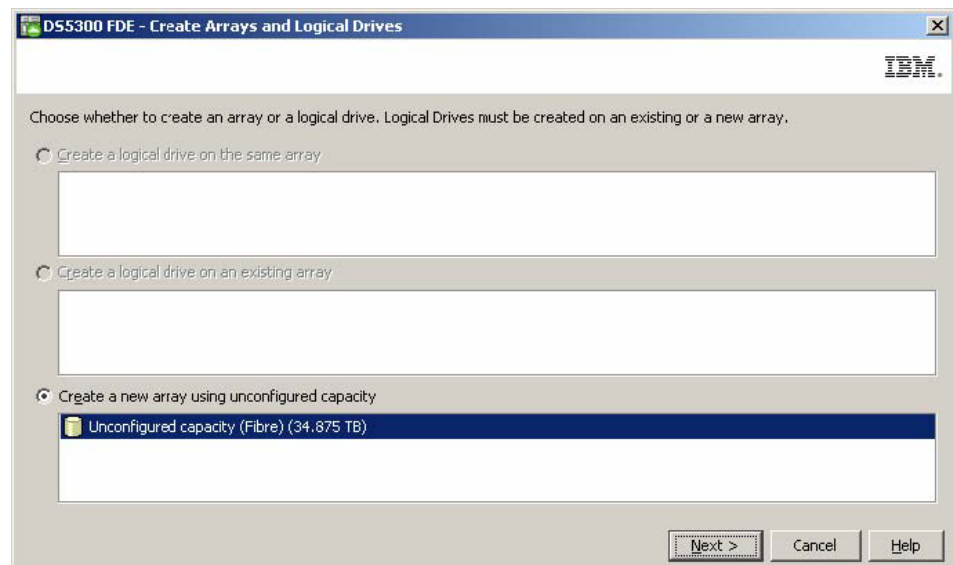
1. Create a RAID array from the FDE drives that are available in the storage subsystem and then secure it. From the Setup page, click **Configure Storage Subsystem**.




2. In the Select Configuration Task window, click **Manual (advanced)**, click **Create arrays and logical drives**, and then click **OK**.



3. In the Create Arrays and Logical Drives window, select **Create a new array using unconfigured capacity**. If other (non-FDE) drive types are also installed in the DS5000, be sure to select only Fibre Channel FDE drives. Click **Next** to continue.

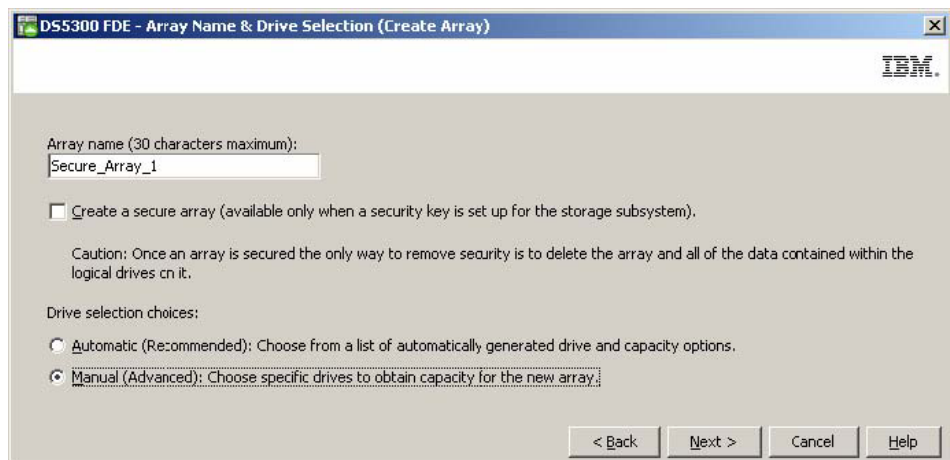


4. Use the Create Array wizard to create the array. Click **Next** to continue.



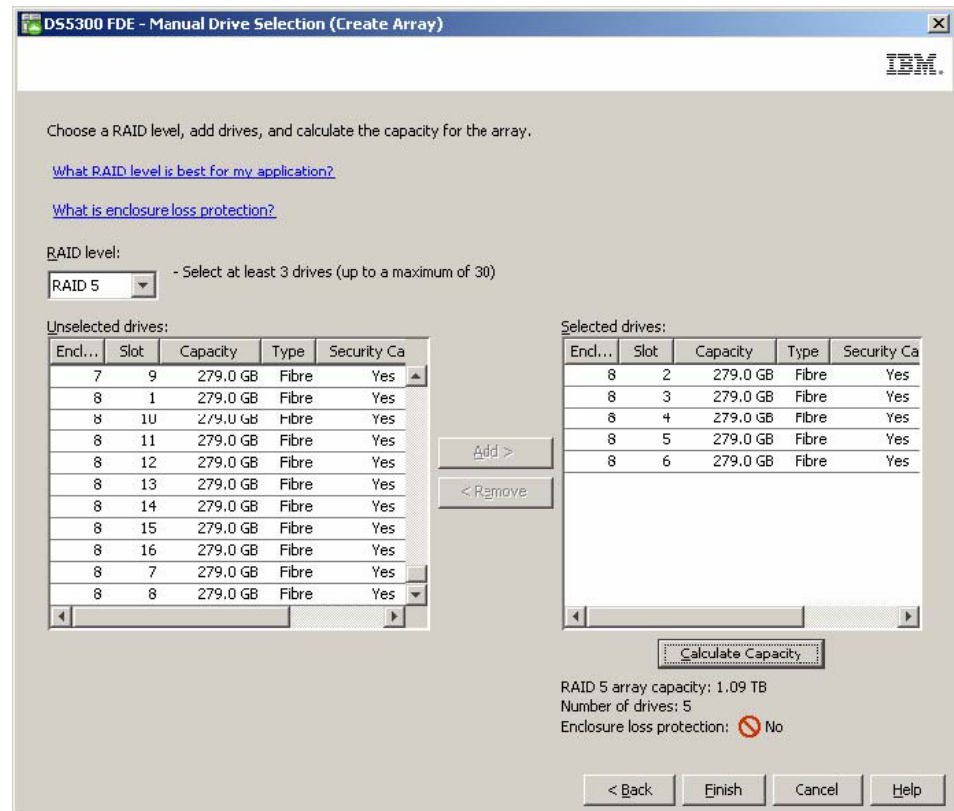
5. In the Array Name & Drive Selection window, enter an array name (for example, Secure_Array_1). Note that the **Create a secure array** check box has been preselected in this window. Clear the **Create a secure array** check box and select **Manual (Advanced)** under **Disk selection choices**. Click **Next** to continue.

Note: The **Create a secure array** check box is displayed and selected *only* if the full disk encryption premium feature is enabled. If you select this check box when you create an array, the array that is created will be secured, and the **Manual (Advanced)** option is not needed to secure the array.

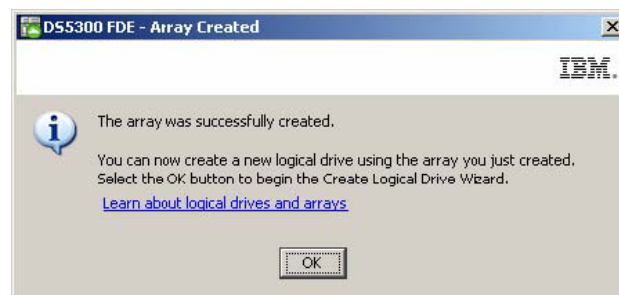


6. Configure drives for the array in the Manual Drive Selection window:
 - a. Select a RAID level (for example, RAID 5).
 - b. From the **Unselected drives** list, select the security-capable drives that you want to use and click **Add** to add them to the **Selected drives** list (for example, select the disk drives in slots 2 through 6 from storage expansion enclosure 8).
 - c. Click **Calculate Capacity** to calculate the total capacity of the selected drives.
 - d. Click **Finish** to complete the array.

Note: These drives are not yet secure. They are secured later in the process.

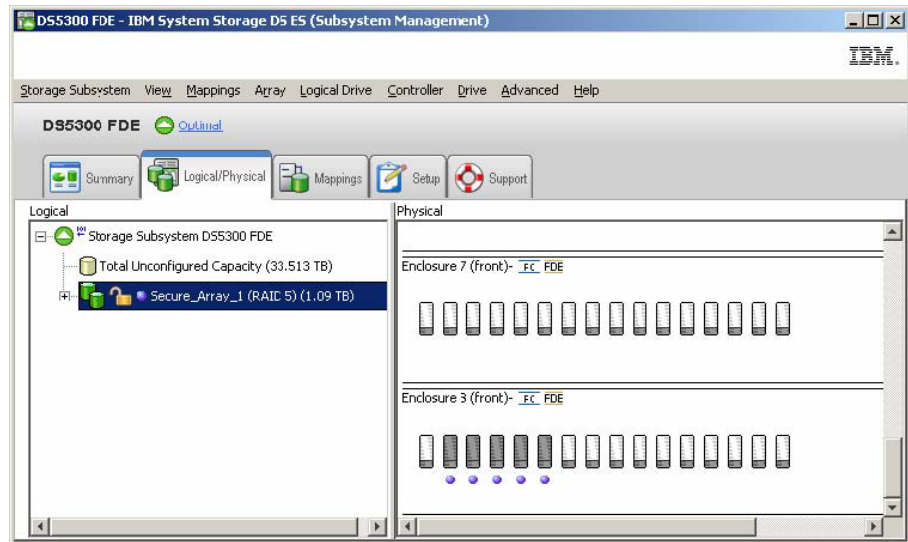


7. In the Array Created window, click **OK** to acknowledge successful creation of the array.

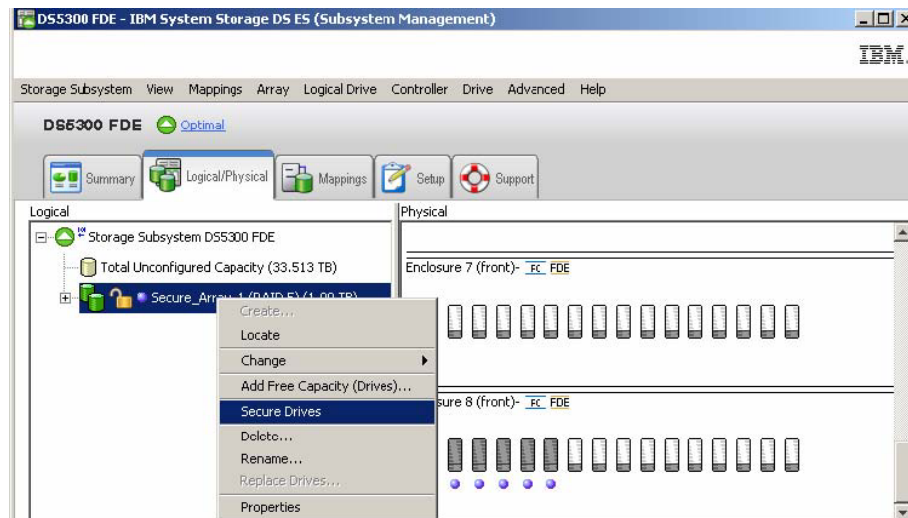


8. When the wizard prompts you to create logical drives in the array, use the wizard to create the logical drives. After the logical drives are created, continue to the next step. See Chapter 4, "Configuring storage," on page 55 for more information about creating logical drives.
9. Secure the array that you have created:
 - a. In the Subsystem Management window, click the **Logical/Physical** tab.


Note: The blue dots below the disk icons on the right side of the window indicate which disks compose the array.



- b. To enable security on the array, right-click the array name; then, click **Secure Drives**.



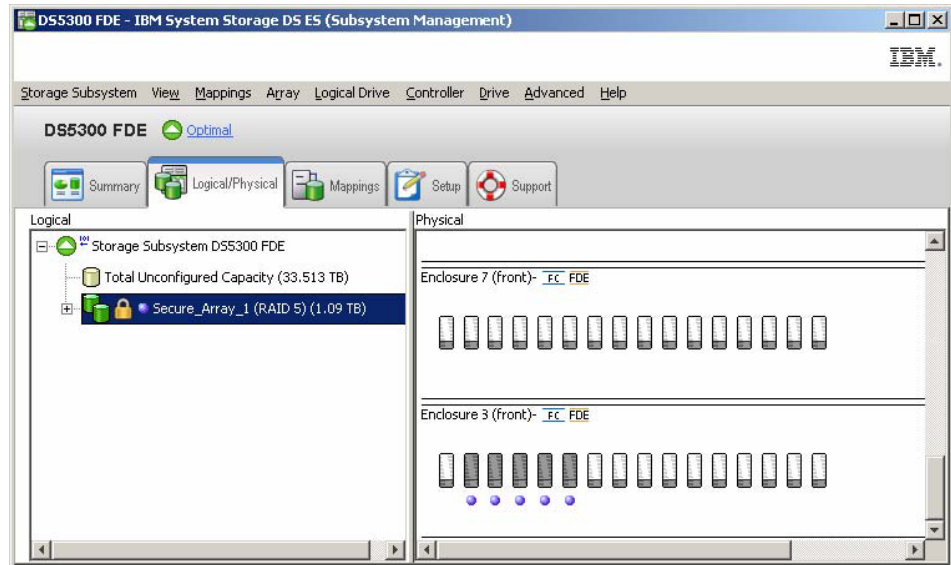
- c. In the Confirm Array Drive Security window, click **Yes** to secure the array.



Note:

- 1) If you move a drive to a separate storage subsystem or if you change the security key more than two times in the current storage subsystem while the drive is removed from the storage subsystem, you must

- provide the pass phrase, the security key, and the security key file to unlock the drive and make the data readable.
- 2) After an array is secured, the only way to remove security is to delete the array. You can make a VolumeCopy of the array and save it to other disks so that the data can continue to be accessed.
10. In the Subsystem Management window, click the **Logical/Physical** tab, and note that the array is secured, as indicated by the lock symbol to the left of the array name.



Unlocking disk drives

A security-enabled FDE drive becomes locked when its power is turned off or when it is removed from the storage subsystem. This is an important feature of storage subsystem disk encryption and FDE drives; the locked state makes the data unreadable to unauthorized users.

Important: If the storage subsystem is in external key management mode and there is not an optimal non-FDE or unsecured FDE drive in the subsystem configuration, you must provide the backup security file and its associated pass phrase to unlock the drives for the storage subsystem to boot successfully.

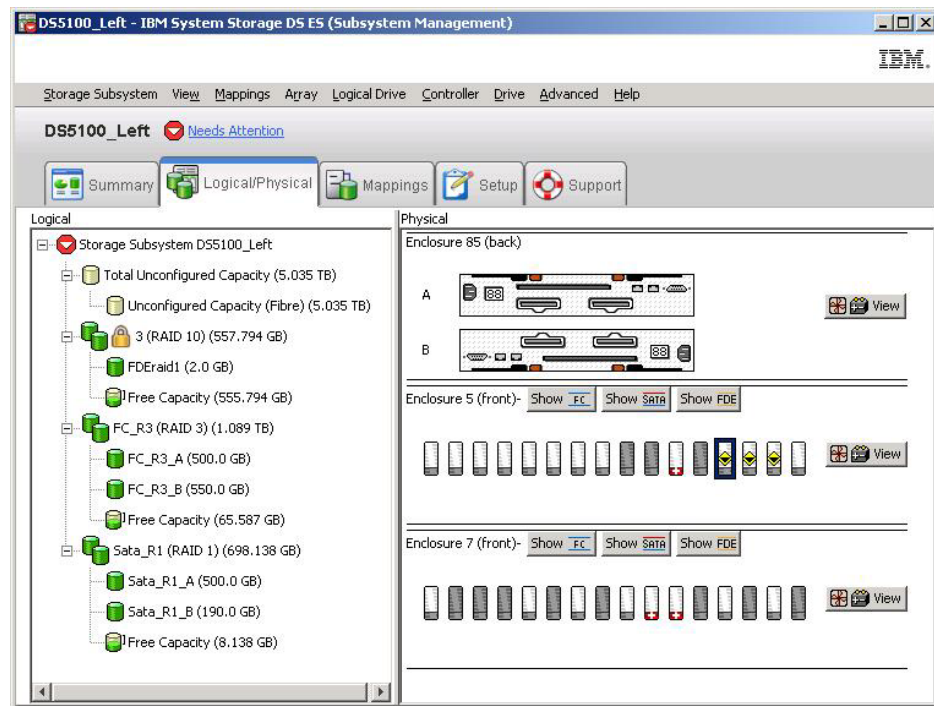
The conditions that cause an FDE drive to become locked vary, depending on the type of security key management that you use. With local security key management, the key is stored inside the controller. Because the controllers always keep a copy of the current and the previous security key, the security key file is not needed every time the storage subsystem power is cycled or a drive is removed and reinserted in the same storage subsystem. However, if a drive is moved to another storage subsystem, or if the security key in the same storage subsystem is changed more than two times while the drive is removed from the storage subsystem, the pass phrase and security file are required to unlock the drive.

Note: Security-enabled FDE drives remain unlocked during firmware updates or while components are replaced. The only time these drives are locked is when they are turned off or removed from the storage subsystem.

With external security key management, the external key manager application supplies the security key to unlock a drive that has been moved from the original subsystem to a new subsystem, provided that the new subsystem is accessible to the application. The new subsystem must be connected to the external key manager application to unlock the drive that was moved. If communication between the external key manager application and the storage subsystem is disrupted, the drives cannot be unlocked until communication is re-established, or until the drives are unlocked with the backup security key file.

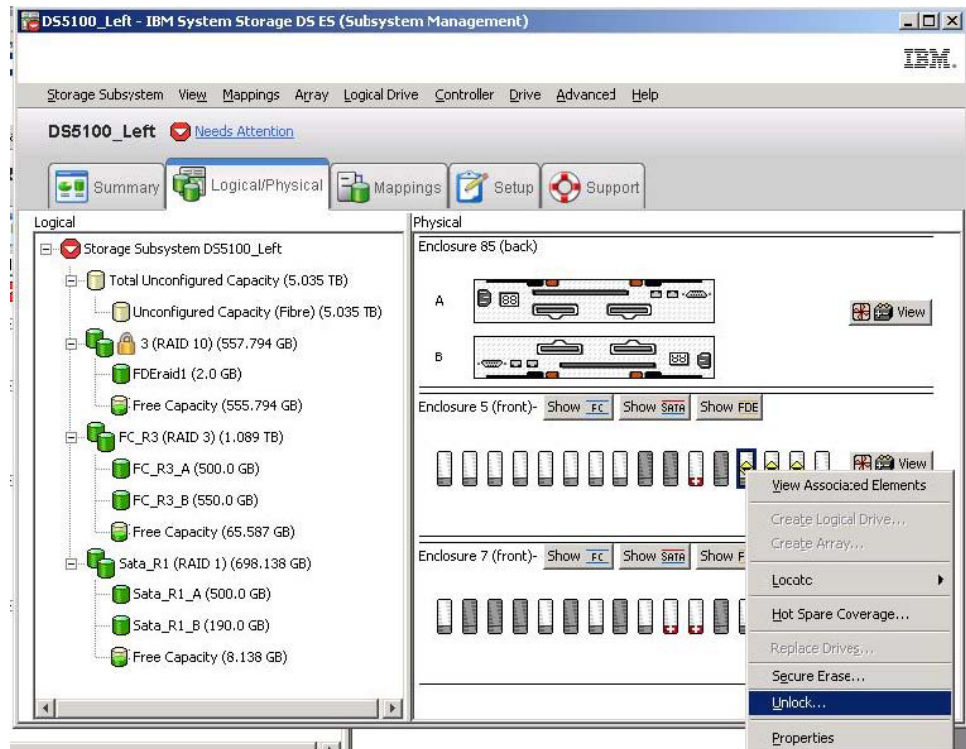
To unlock a locked FDE drive with the backup security key file, complete the following steps:

1. In the Subsystem Management window, click the **Logical/Physical** tab.



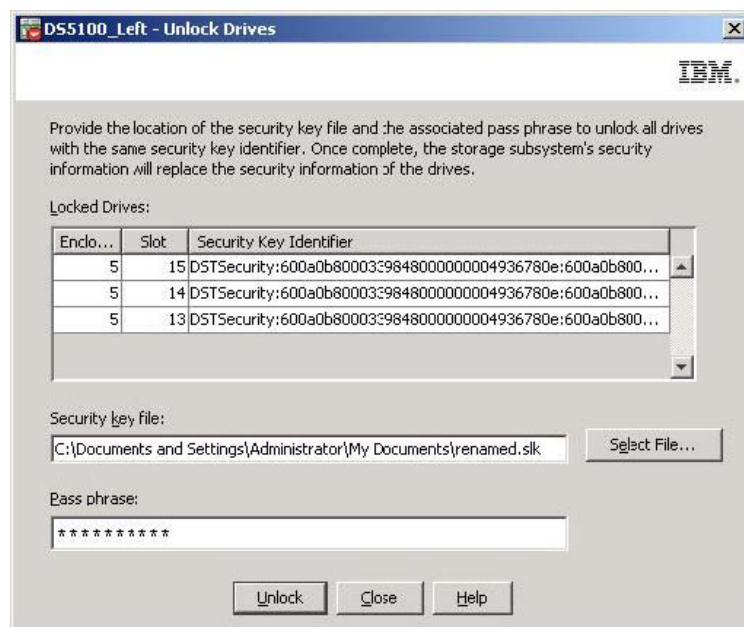
2. Right-click the drives that you want to unlock; then, click **Unlock**.

Note: If you want to unlock multiple drives, you only have to select one drive. The Storage Manager automatically lists all of the drives that are locked in the storage subsystem and checks each drive against the supplied security key file to determine whether it can use the key in the security key file.



3. In the Unlock Drives window, the locked drives that you selected are listed. To unlock these drives, select the security key file, enter the pass phrase, and then click **Unlock**. The storage subsystem uses the pass phrase to decrypt the security key from the security key file. The storage subsystem then compares the decrypted security key to the security key on the drive and unlocks all the drives for which the security key matches.

Note: The authentication process occurs only when the drive is in Locked state because the drive was powered on after a power-down event. It does not repeat with each read and write operation.



4. In the Unlock Drives Complete window, click **OK** to confirm that the drives are unlocked. The unlocked drives are now ready to be imported.




Migrating storage subsystems (head-swap) with FDE drives

With an FDE-compatible storage subsystem, you can migrate drives as a complete storage subsystem into another FDE-compatible storage subsystem with existing disk group migration techniques. User data remains intact on the disks because configuration metadata is stored on every drive in the storage subsystem. FDE security-enabled drives can also be migrated and remain secure with a few additional steps that are described in this section.

Note:

1. The following procedure describes only the additional data migration steps that are required for secure arrays. For complete information and procedures, see the *IBM System Storage DS3000, DS4000, or DS5000 Hard Drive and Storage Expansion Enclosure Installation and Migration Guide*.
2. The following data migration steps also apply when you replace both controllers in the storage subsystem. All drives in that storage subsystem must be included. Partial migrations are not supported when you replace both controllers. A security file is necessary in this case; you might not have management access to the storage subsystem to export the current security key if both of the controllers must be replaced.
1. Save the security key that is used to unlock the drives in the existing storage subsystem in a security key file before you remove the drives from the existing storage subsystem. After you export the security key, pass phrase, and security key file, the security key file can be transferred from one storage subsystem to another.
 - a. In the Subsystem Management window, click **Storage Subsystem**, click **Drive Security**, and click **Save Security Key File**.

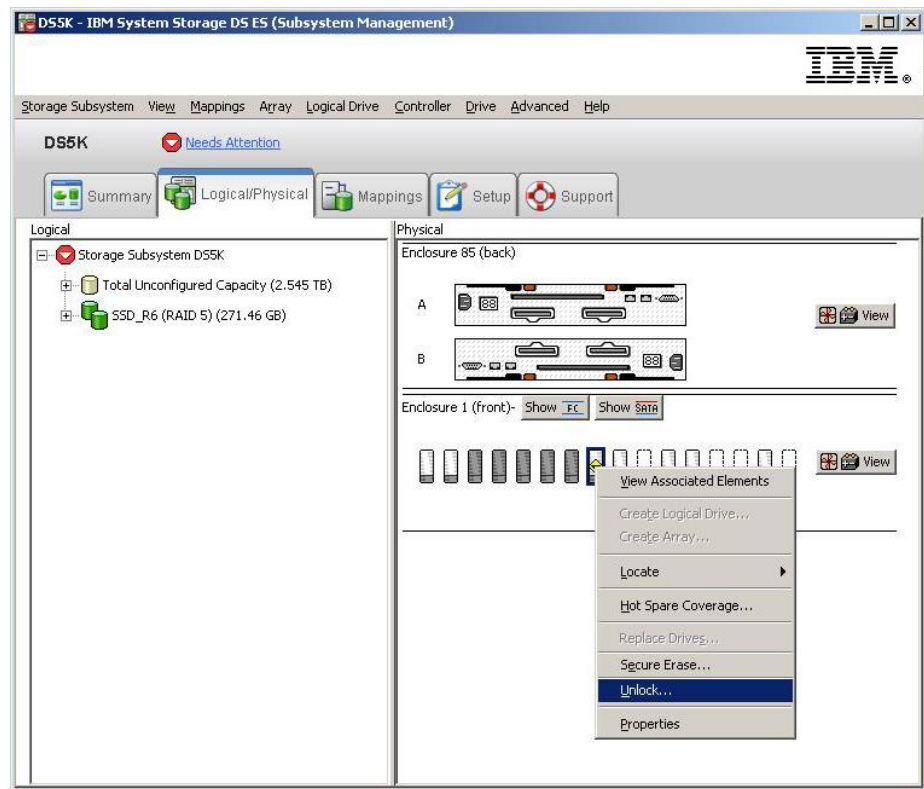


- b. In the Save Security Key File - Enter Pass Phrase window, select a file save location, and enter and confirm the pass phrase; then, click **Save**.

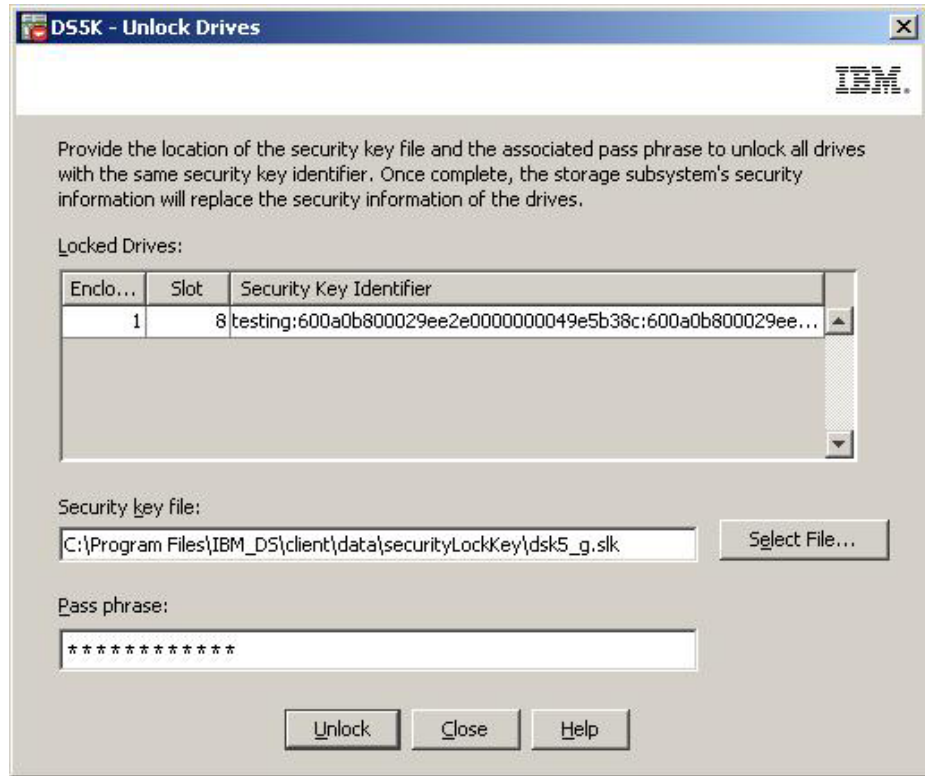


- c. Export the defined arrays in the original storage subsystem.
 - d. Turn off the subsystem power and replace the old storage subsystem controller enclosure with the new controller enclosure.
 - e. Turn on the power to the new storage subsystem.
2. After you replace the existing storage subsystem controller enclosure with a new controller enclosure, unlock the security-enabled FDE drives before you import the RAID arrays:
 - a. Click the **Logical/Physical** tab in the Subsystem Management window.
 - b. Right-click the drives that you want to unlock; then, click **Unlock**.

Note: The Full Disk Encryption premium feature might be incompliant. Generate the new FDE premium feature key file to enable the storage subsystem FDE functionality.



- c. Select the security key file for the selected drives and enter the pass phrase that you entered when saving the security key back up file; then, click **Unlock**.



Erasing disk drives

Attention: All data on the disk will be permanently erased when the secure-erase operation is completed on a security-enabled FDE drive. Do not perform this action unless you are sure that you want to erase the data.

Secure erase provides a higher level of data erasure than other traditional methods. When you initiate secure erase with the Storage Manager, a command is sent to the FDE drive to perform a cryptographic erase. A cryptographic erase erases the existing data encryption key and then generates a new encryption key inside the drive, making it impossible to decrypt the data. After the encryption key is changed, any data that was written to the disk that was encrypted with the previous encryption key is unintelligible. This includes all bits, headers, and directories.


After secure erase takes place on the drive, the following actions occur:

- The data becomes completely and permanently inaccessible, and the drive returns to the original factory state.
- Drive security becomes disabled and must be re-enabled if it is required.

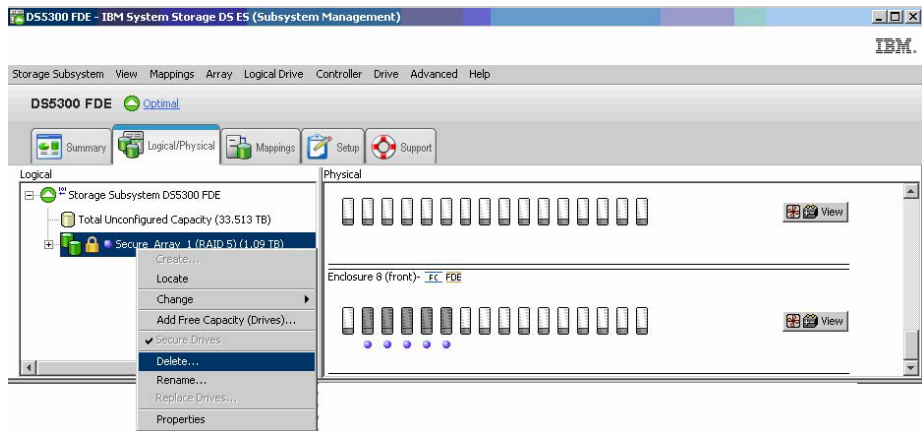
Before you initiate secure erase, the security-enabled FDE drive must be unlocked, and the array that it is assigned to must be deleted.

Attention: You must back up the data in the security-enabled FDE drives to other drives or to secure tape before you secure erase an FDE drive if you want to access the data at a later time. All data on the disk will be permanently erased when the secure-erase operation is completed on a security-enabled FDE drive. Do not perform this action unless you are sure that you want to erase the data. The improper use of secure erase will result in lost data.

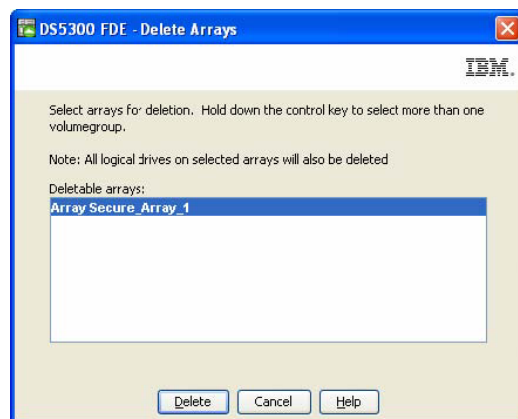
1. Before the drives can be secure erased, you must delete the RAID array that the drives are associated with and return the drives to Unassigned status:
 - a. Click the **Logical/Physical** tab in the Subsystem Management window.



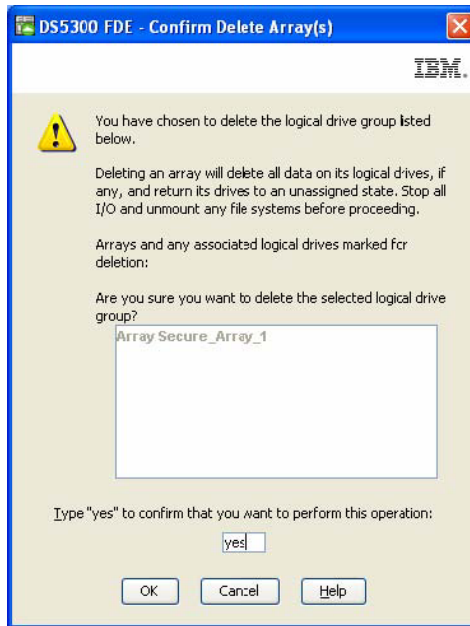
- b. Right-click the array name; then, click **Delete**.




- c. When you are prompted to select the array that you want to delete, click the array name and click **Delete**.



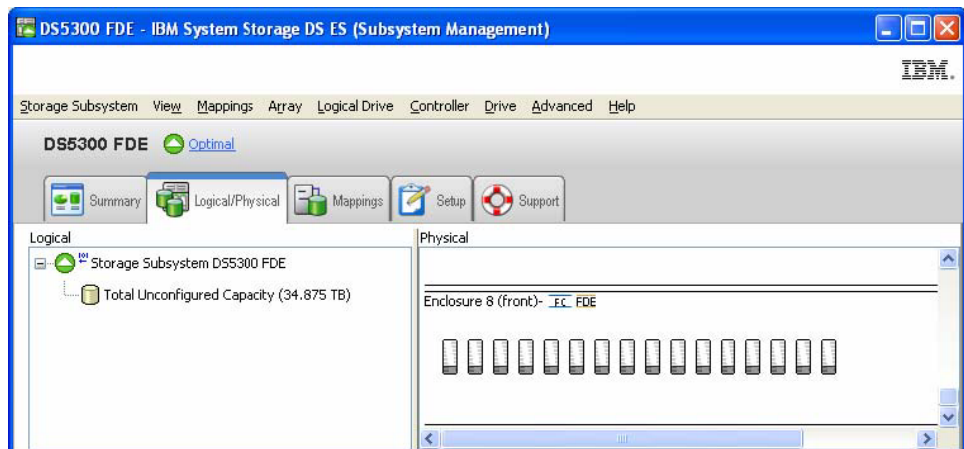
- d. To confirm that you want to delete the array, enter yes in the field and click **OK**.




- e. Wait for the array deletion process to be completed. When you receive the confirmation Processed 1 of array(s) – Complete, click **OK**.




2. Click the **Logical/Physical** tab in the Subsystem Management window.



3. Select the drive on which you want to perform a secure erase. You can select more than one drive to be erased by holding down the Ctrl key. In the top menu bar, click **Drive**; then, click **Secure Erase**.



4. To confirm that you want to permanently erase all data on the disk, enter yes in the field and click **OK**. These drives can now be repurposed or discarded.



Global hot-spare disk drives

If a disk drive fails in a FDE-compatible storage subsystem, the controller uses redundant data to reconstruct the data on the failed drive on a global hot-spare drive. The global hot-spare drive is automatically substituted for the failed drive without intervention. When the failed drive is eventually replaced, the data from the hot-spare drive is copied back to the replacement drive.

Hot-spare drives must meet the array hot-spare requirements. The following drive types are required for hot-spare drives when secure-capable arrays are configured. If a drive does fail, the Storage Manager automatically determines which hot-spare drive to substitute according to the type of the failed drive.

- For an array that has secured FDE drives, the hot-spare drive must be an unsecured FDE drive of the same or greater capacity. After the unsecured FDE hot-spare drive is used as a spare for a failed drive in the secured RAID array, it is security enabled.
- In an array that has FDE drives that are not secured, the hot-spare drive can be either an unsecured FDE drive or a non-FDE drive.

Note: If an unsecured FDE hot-spare drive was used as a spare for a non-secured FDE array and the array was secured after the data was copied back, the

unsecured FDE hot-spare drive remains unsecured, exposing the data in the drive if it is removed from the storage subsystem.

An unconfigured secured FDE drive cannot be used as a global hot-spare drive. If a global hot spare is a secured FDE drive, it can be used as a spare drive only in secured arrays. If a global hot-spare drive is an unsecured FDE drive, it can be used as a spare drive in secured or unsecured arrays with FDE drives, or as a spare drive in arrays with non-FDE drives. You must secure erase the FDE drive to change it to Unsecured state before it can be used as a global hot-spare drive. The following error message is generated if you assign an unconfigured secured FDE drive as a global hot spare.

Return code: Error 2 - The operation cannot complete because either (1) the current state of a component does not allow the operation to be completed, (2) the operation has been disabled in NVSRAM (example, you are modifying media scan parameters when that option (offset 0x31, bit 5) is disabled), or (3) there is a problem with the storage subsystem. Please check your storage subsystem and its various components for possible problems and then retry the operation. Operation when error occurred:
PROC_assignSpecificDrivesAsHotSpares

When a global hot-spare drive is used as a spare for a failed drive in a secure array, it becomes a secure FDE drive and remains secure provided that it is a spare in the secure array. After the failed drive in the secure array is replaced and the data in the global hot-spare drive is copied back to the replaced drive, the global hot-spare drive is automatically reprovisioned by the controllers to become an unsecured FDE global hot-spare drive.

As a best practice in a mixed disk environment that includes non-security-capable SATA drives, non-security-capable Fibre Channel drives, and FDE Fibre Channel drives (with security enabled or not enabled), use at least one type of global hot-spare drive (FDE Fibre Channel and a SATA drive) at the largest capacity within the array. If a secure-capable FDE Fibre Channel and a SATA hot-spare drive are included, all arrays are protected.

Follow the standard hot-spare drive configuration guidelines in “Configuring global hot-spare drives” on page 80). Hot-spare configuration guidelines are the same for FDE drives.

Log files

The Storage Manager major events log (MEL) includes messages that describe any security changes that are made in the storage subsystem.

Frequently asked questions

This section lists frequently asked questions about FDE. The questions and answers are organized in the following categories:

- “Securing arrays” on page 240
- “Secure erase” on page 240
- “Local security key management” on page 241
- “External security key management” on page 241
- “Premium features” on page 241
- “Global hot-spare drives” on page 242
- “Boot support” on page 242

- “Locked and unlocked states” on page 242
- “Backup and recovery” on page 242
- “Other” on page 243

Securing arrays

- Can I change an unsecured array with FDE drives to a secured array?
 - Yes. The steps to complete this process are described in “Securing a RAID array” on page 224. The DS5000 Encryption feature must be enabled and the security key file and pass phrase already established. See “Enabling premium features” on page 217 for more information.
- When I enable security on an array, will the data that was previously written to that array be lost or erased?
 - No. Unless you perform a secure erase on the array disk drives, this data remains intact.
- Can I change a secured array with FDE drives to an unsecured array?
 - No. This is not a supported option. After an unsecured array is changed to a secure array, you cannot change it back to an unsecured array without destroying the data in the security-enabled FDE drives. Use VolumeCopy to copy the secure data to an unsecured array, or back up the data to a secured tape. If you VolumeCopy the secure data to an unsecured array, you must physically secure the drives. Then you must delete the original array and secure erase the array drives. Create a new unsecured array with these drives and use VolumeCopy to copy the data back to the original drives, or restore the data from secure tape.
- If I have an array with FDE drives that is secured, can I create another array that uses these same drives and not enable security? Does the storage subsystem have a control so that this does not occur?
 - No. These are not supported functions. Any logical drive that is part of an array must be secured, because the drives on which it is stored are security enabled.
- When a secure array is deleted, does disk security remain enabled?
 - Yes. The only way to disable security is to perform a secure erase or reprovision the drives.
- If I create a new array on a set of unassigned/unconfigured security-enabled FDE disks, will they automatically become secure?
 - Yes.

Secure erase

- With secure erase, what can I erase, an individual drive or an array?
 - Secure erase is performed on an individual drive. You cannot erase a secured drive that is part of an array; you must first delete the array. After the array is deleted and the drives become unassigned, you can erase multiple disks in the same operation by holding the Ctrl key while you select the drives that are to be secure erased.
- If I want to use only the secure erase feature, do I still have to set up a security key identifier and pass phrase?
 - Yes. The full disk encryption feature must be enabled before you can use secure erase.
- After secure erase is performed on a drive, is security enabled or disabled on that drive?

- The drive is returned to Secure Capable (unsecured) state after a secure erase. Security is disabled on the drive.
- If I inadvertently secure erase a drive, is it possible to recover the data in the drive?
 - No. After a drive is secure erased, the data in the drive is not recoverable. You must recover the lost data from a backup copy. Back up the data in secure drives before you secure erase the drives.

Local security key management

- Can I get to the security keys through the Storage Manager or controller?
 - No. The security key is obfuscated in the storage subsystem. Only an encrypted version of the key can be exported to a security key file, using the save security key operation. The actual security key is not available for viewing. Implement prudent security features for the storage subsystem. The Storage Manager forces a strong password, but administrator access must have stringent controls in place.
- If I lose a drive that is unlocked or security disabled, can that data be accessed even though the data is encrypted?
 - Yes. Because security has not been enabled on the drive, it remains unlocked, and the data is accessible.
- If my security key falls into the wrong hands, can I change it without losing my data?
 - Yes. The drive can be re-keyed, using the procedure to change the security key.

External security key management

- How is external security key management different from local security key management?
 - Instead of using a security key that is housed and obfuscated in a storage subsystem controller, external security key management uses a central key location on your network to manage keys for different storage subsystems. External security key management is facilitated by external key license manager software, such as IBM Tivoli Key Lifecycle Manager (TKLM). If you do not already have this software, you must purchase it, install it, and configure the proxy server to set up external security key management.
- Do I need access to the saved security file when I move secured drives from one storage subsystem to another?
 - No. If the new storage subsystem is connected and recognized by the proxy server and the external key management software, the software provides the security key to unlock the drive automatically.
- Why does the storage subsystem require that I manually supply the security key from the saved security file after I cycled the power to the subsystem?
 - The subsystem does not have at least one installed non-FDE drive.

Premium features

- How do I make sure that my mirrored data is secure? What is a best practice for protecting data at the remote site?
 - Secure your data with security-enabled FDE drives at both the primary and secondary sites. Also, you must make sure that the data is protected while it is being transferred between the primary and secondary sites.

- Can I use VolumeCopy to copy a secured logical unit number to a unsecured one? If so, what prevents someone from doing that first and then stealing the unsecured copy?
 - Yes. To prevent someone from stealing the data with this method, implement prudent security features for the DS5000 storage subsystem. The Storage Manager forces a strong password, but administrator access must also have stringent controls in place.
- Can FlashCopy and VolumeCopy data be secured?
 - Yes. For FlashCopy, the FlashCopy repository logical drive must be secured if the target FlashCopy data is secured. The Storage Manager enforces this rule. Similarly, if the source array of the VolumeCopy pair is secured, the target array of the VolumeCopy pair must also be secured.

Global hot-spare drives

- Can I use an unconfigured FDE drive as a global hot-spare drive?
 - Yes, but only if the drive is unsecured (security not enabled). Check the status of the unconfigured FDE drive. If the drive is secure, it must be secure erased or reprovisioned before you can use it as a global hot-spare drive.
- If the hot-spare drive in a secured array is an unsecured FDE drive, does this drive automatically become secured when a secured FDE drive fails and that data is written to the hot-spare drive?
 - Yes. When the failed drive is removed from the RAID group, a rebuild is automatically started to the hot-spare drive. Security is enabled on the hot-spare drive before the rebuild is started. A rebuild cannot be started to a non-FDE drive for a secure array. After the failed drive in the secured array is replaced and the data in the global hot-spare drive is copied back to the replaced drive, the global hot-spare drive is automatically reprovisioned by the controllers to become an unsecured FDE global hot-spare drive.

Boot support

- Is there a special process for booting from a security-enabled drive?
 - No. The only requirement is that the storage subsystem must be running (which is required in any booting process).
- Are FDE drives susceptible to cold boot attacks?
 - No. This issue applies more to the server side, because an individual can create a boot image to gain access to the server. This does not apply to FDE drives. FDE drives do not use the type of memory that is susceptible to a cold boot attack.

Locked and unlocked states

- When does a security-enabled drive go into Locked state?
 - The drive becomes locked whenever the disk is powered off. When the FDE drive is turned off or disconnected, it automatically locks down the data on the disk.

Backup and recovery

- How can I make sure that my archived data is secure?
 - Securing archived data is beyond the scope of this document. See the Storage Networking Interface Association (SNIA) guidelines for secure tape backup. For specific references, see the *IBM Full Disk Encryption Best Practices* document. To access this document on the IBM website, go to

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5081492&brandind=5000028>, or complete the following steps:

1. Go to the IBM Support Portal at <http://www.ibm.com/support/entry/portal>.
2. In the **Search within all of support & downloads** field at the bottom of the webpage, type FDE and press Enter.
3. In the list of search results, click the **IBM Full Disk Encryption Best Practices - IBM System Storage** link.
4. Click the link to the PDF file to open or download the *IBM Full Disk Encryption Best Practices* document.

Other

- Is DACstore information still written to the disk?
 - Yes. However, if the drive is secured, it must be unlocked by the controller before the DACstore information can be read. In the rare event that the controller security key is corrupted or both controllers are replaced, a security key file must be used to unlock the drive.
- Is data on the controllers cache secure with FDE and IBM Disk Encryption? If not, are there any best practices here?
 - No. This is a security issue of physical access to the hardware. The administrator must have physical control and security of the storage subsystem itself.
- If I have secure-capable disks but have not purchased the IBM Disk Encryption premium feature key, can I still recognize secure-capable disks from the user interface?
 - Yes. This information is available from several windows in the Storage Manager interface.
- What about data classification?
 - See the SNIA best practices for more information about data classification. For specific references, see the *IBM Full Disk Encryption Best Practices* document. To access this document on the IBM website, go to <http://www-947.ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5081492&brandind=5000028>, or complete the following steps:
 1. Go to the IBM Support Portal at <http://www.ibm.com/support/entry/portal>.
 2. In the **Search within all of support & downloads** field at the bottom of the webpage, type FDE and press Enter.
 3. In the list of search results, click the **IBM Full Disk Encryption Best Practices - IBM System Storage** link.
 4. Click the link to the PDF file to open or download the *IBM Full Disk Encryption Best Practices* document.
- Can I use both FDE and non-FDE drives if I do not secure the drives?
 - Yes. However, using both FDE and non-FDE drives is not a cost-effective use of FDE drives. An array with both FDE and non-FDE drives cannot be converted into a secure array at a later time.
- Do FDE disk drives have lower usable capacity because the data is encrypted or because capacity is needed for the encryption engine and keys?
 - No. There is no capacity difference between non-FDE and FDE disk drives (1 GB unencrypted = 1 GB encrypted).

Chapter 7. Troubleshooting

Use the information in this chapter to diagnose and solve problems related to Storage Manager. For information about getting help, service, or other technical assistance, see “Getting information, help, and service” on page xvi.

The following topics are covered in this chapter:

- “Critical event problem solving”
-
- “DS Diagnostic Data Capture (DDC)” on page 263
- “Resolving disk array errors on AIX” on page 266

Critical event problem solving

When a critical event occurs, it is logged in the event log and sent to any email and SNMP trap destinations that are configured. The critical event type and the sense key, ASC, and ASCQ data are shown in the event log details.

If a critical event occurs and you plan to call IBM support, you can use the Customer Support Bundle feature to gather and package various pieces of data that can aid in remote troubleshooting. To use the Customer Support Bundle feature, complete the following steps:

1. From the Subsystem Management window of the logical drive that is exhibiting problems, click **Advanced > Troubleshooting > Advanced > Collect All Support Data**. The Collect All Support Data window opens.
2. Type the name of the file where you want to save the collected data or browse to select the file. Click **Start**.

Note: It takes several minutes for the compressed file to be created, depending on the amount of data that is to be collected.

3. When the process is complete, you can send the compressed file electronically to IBM support for troubleshooting.

Table 38 provides more information about events with a critical priority, as shown in the Subsystem Management window event log.

Table 38. Critical events

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 1001 - Channel failed	6/3F/C3	<p>Description: The controller failed a channel and cannot access drives on this channel anymore. The FRU group qualifier (byte 26) in the sense data indicates the relative channel number of the failed channel. Typically, this condition is caused by a drive that is ignoring the SCSI protocol on one of the controller destination channels. The controller fails a channel if it issued a reset on a channel and continues to see the drives ignore the SCSI Bus Reset on this channel.</p> <p>Action: Start the Recovery Guru to access the Failed Drive SCSI Channel recovery procedure. Contact IBM support to complete this procedure.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 1010 - Impending drive failure (PFA) detected	6/5D/80	<p>Description: A drive has reported that a failure prediction threshold has been exceeded. This indicates that the drive might fail within 24 hours.</p> <p>Action: Start the Recovery Guru and click the Impending Drive Failure recovery procedure. Follow the instructions to correct the failure.</p>
Event 1015 - Incorrect mode parameters set on drive	6/3F/BD	<p>Description: The controller is unable to query the drive for its current critical mode page settings or is unable to change these settings to the correct setting. This indicates that the Qerr bit is set incorrectly on the drive specified in the FRU field of the Request Sense data.</p> <p>Action: The controller has not failed yet. Contact IBM support for the instructions to recover from this critical event.</p>
Event 1207 - Fibre Channel link errors - threshold exceeded	None	<p>Description: Invalid characters have been detected in the Fibre Channel signal. Possible causes for the error are a degraded laser in a gigabit interface converter (GBIC) or media interface adapter, damaged or faulty Fibre Channel cables, or poor cable connections between components on the loop.</p> <p>Action: In the main Subsystem Management window, click Help → Recovery Procedures. Click Fibre Channel Link Errors Threshold Exceeded for more information about recovering from this failure.</p>
Event 1208 - Data rate negotiation failed	None	<p>Description: The controller cannot auto-negotiate the transfer link rates. The controller considers the link to be down until negotiation is attempted at controller start-of-day or when a signal is detected after a loss of signal.</p> <p>Action: Start the Recovery Guru to access the Data Rate Negotiation Failed recovery procedure and follow the instructions to correct the failure.</p>
Event 1209 - Drive channel set to Degraded	None	<p>Description: A drive channel status was set to Degraded because of excessive I/O errors or because a technical support representative advised the array administrator to manually set the drive channel status for diagnostic or other support reasons.</p> <p>Action: Start the Recovery Guru to access the Degraded Drive Channel recovery procedure and follow the instructions to correct the failure.</p>
Event 150E - Controller loopback diagnostics failed	None	<p>Description: The controller cannot initialize the drive-side Fibre Channel loops. A diagnostic routine has been run and has identified a controller problem, and the controller has been placed offline. This event occurs only on certain controller models.</p> <p>Action: Start the Recovery Guru to access the Offline Controller recovery procedure and follow the instructions to replace the controller.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 150F - Channel miswire	None	<p>Description: Two or more drive channels are connected to the same Fibre Channel loop. This can cause the storage subsystem to behave unpredictably.</p> <p>Action: Start the Recovery Guru to access the Channel Miswire recovery procedure and follow the instructions to correct the failure.</p>
Event 1510 - ESM blade miswire	None	<p>Description: Two ESM blades in the same storage expansion enclosure are connected to the same Fibre Channel loop. A level of redundancy has been lost, and the I/O performance for this storage expansion enclosure is reduced.</p> <p>Action: Start the Recovery Guru to access the ESM blade Miswire recovery procedure and follow the instructions to correct the failure.</p>
Event 1513 - Individual Drive - Degraded Path	None	<p>Description: The specified drive channel is experiencing intermittent errors along the path to a single drive or to several drives.</p> <p>Action: Start the Recovery Guru to access the Individual Drive - Degraded Path recovery procedure and follow the instructions to recover from this failure.</p>
Event 1600 - Uncertified drive detected	None	<p>Description: An uncertified drive has been inserted into the storage subsystem.</p> <p>Action: Start the Recovery Guru to access the Uncertified Drive recovery procedure and follow the instructions to recover from this failure.</p>
Event 1601 - Reserved blocks on ATA drives cannot be discovered	None	<p>Description: Reserved blocks on the ATA drives are not recognized.</p> <p>Action: Contact IBM support for instructions for recovering from this event.</p>
Event 200A - Data/parity mismatch detected on logical drive	None	<p>Description: A media scan operation has detected inconsistencies between a portion of the data blocks on the logical drive and the associated parity blocks. User data in this portion of the logical drive might have been lost.</p> <p>Action: Select an application-specific tool (if available) to verify that the data is correct on the logical drive. If no such tool is available, or if problems with the user data are reported, restore the entire logical drive contents from the most recent backup, if the data is critical.</p>
Event 202E - Read drive error during interrupted write	3/11/8A	<p>Description: A media error has occurred on a read operation during interrupted write processing.</p> <p>Action: Start the Recovery Guru to access the Unrecovered Interrupted Write recovery procedure. Contact IBM support to complete this procedure.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2109 - Controller cache not enabled - cache sizes do not match	6/A1/00	<p>Description: The controller cannot enable mirroring if the alternative controller cache size of both controllers is not the same. Verify that the cache size for both controllers is the same.</p> <p>Action: Contact IBM support for the instructions for recovering from this failure.</p>
Event 210C - Controller cache battery failed	6/0C/80	<p>Description: The controller has detected that the battery is not physically present, is fully discharged, or has reached its expiration date.</p> <p>Action: Start the Recovery Guru to access the Failed Battery CRU recovery procedure and follow the instructions to correct the failure.</p>
Event 210E - Controller cache memory recovery failed after power cycle or reset	6/0C/81	<p>Description: Recovery from a data-cache error was unsuccessful. User data might have been lost.</p> <p>Action: Contact IBM support for the instructions for recovering from this failure.</p>
Event 2110 - Controller cache memory initialization failed	6/40/81	<p>Description: The controller has detected the failure of an internal controller component (RAID buffer). The internal controller component failure might have been detected during operation or during an on-board diagnostic routine.</p> <p>Action: Contact IBM support for the instructions for recovering from this failure.</p>
Event 2113 - Controller cache battery nearing expiration	6/3F/D9	<p>Description: The cache battery is within six weeks of its expiration.</p> <p>Action: Start the Recovery Guru to access the Battery Nearing Expiration recovery procedure and follow the instructions to correct the failure.</p>
Event 211B - Batteries present but NVSRAM configured for no batteries	None	<p>Description: A battery is present in the storage subsystem, but the NVSRAM is set to not include batteries.</p> <p>Action: Contact your IBM technical support representative for the instructions to recover from this failure.</p>
Event 2229 - Drive failed by controller	None	<p>Description: The controller failed a drive because of a problem with the drive.</p> <p>Action: Start the Recovery Guru to access the Drive Failed by Controller procedure and follow the instructions to correct the failure.</p>
Event 222D - Drive manually failed	6/3F/87	<p>Description: The drive was manually failed by a user.</p> <p>Action: Start the Recovery Guru to access the Drive Manually Failed procedure and follow the instructions to correct the failure.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2247 - Data lost on the logical drive during unrecovered interrupted write	6/3F/EB	<p>Description: An error has occurred during interrupted write processing during the start-of-day routine, which caused the logical drive to go into a failed state.</p> <p>Action: Start the Recovery Guru to access the Unrecovered Interrupted Write recovery procedure and follow the instructions to correct the failure. Contact IBM support to complete this procedure.</p>
Event 2248 - Drive failed - write failure	6/3F/80	<p>Description: The drive failed during a write command. The drive is marked Failed.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2249 - Drive capacity less than minimum	6/3F/8B	<p>Description: During drive replacement, the capacity of the new drive is not large enough to support all the logical drives that must be reconstructed on it.</p> <p>Action: Replace the drive with a larger capacity drive.</p>
Event 224A - Drive has wrong block size	6/3F/8C	<p>Description: The drive block size does not match that of the other drives in the logical drive. The drive is marked Failed.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 224B - Drive failed - initialization failure	6/3F/86	<p>Description: The drive failed from either a Format Unit command or a Write operation (issued when a logical drive was initialized). The drive is marked Failed.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 224D - Drive failed - no response at start of day	6/3F/85	<p>Description: The drive failed a Read Capacity or Read command during the start-of-day routine. The controller is unable to read the configuration information that is stored on the drive. The drive is marked Failed.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 224E - Drive failed - initialization/reconstruction failure	6/3F/82	<p>Description: The previously failed drive is marked Failed because of one of the following reasons:</p> <ul style="list-style-type: none"> • The drive failed a Format Unit command that was issued to it • The reconstruction on the drive failed because the controller was unable to restore it (for example, because of an error that occurred on another drive that was required for reconstruction). <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2250 - Logical drive failure	6/3F/E0	<p>Description: The controller has marked the logical drive Failed. User data and redundancy (parity) can no longer be maintained. The most likely cause is the failure of a single drive in nonredundant configurations or a nonredundant second drive in a configuration that is protected by one drive.</p> <p>Action: Start the Recovery Guru to access the Failed Logical Drive Failure recovery procedure and follow the instructions to correct the failure.</p>
Event 2251 - Drive failed - reconstruction failure	6/3F/8E	<p>Description: A drive failed because of a reconstruction failure during the start-of-day routine.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2252 - Drive marked offline during interrupted write	6/3F/98	<p>Description: An error has occurred during interrupted write processing, which caused the logical drive to be marked Failed. Drives in the array that did not experience the read error go into the Offline state and log this error.</p> <p>Action: Start the Recovery Guru to access the Unrecovered Interrupted Write recovery procedure. Contact IBM support to complete this procedure.</p>
Event 2254 - Redundancy (parity) and data mismatch is detected	6/8E/01	<p>Description: The controller detected inconsistent redundancy (parity) or data during a parity verification.</p> <p>Action: Contact IBM support for the instructions for recovering from this failure.</p>
Event 2255 - Logical drive definition incompatible with ALT mode - ALT disabled Note: This event is not applicable for the DS4800.	6/91/3B	<p>Description: Auto-LUN transfer (ALT) works only with arrays in which only one logical drive is defined. Currently, there are arrays in the storage subsystem in which more than one logical drive is defined; therefore, ALT mode has been disabled. The controller operates in normal redundant controller mode, and if there is a problem, it transfers all logical drives on an array instead of transferring individual logical drives.</p> <p>Action: Contact IBM support for the instructions for recovering from this failure.</p>
Event 2260 - Uncertified drive	ASC/ASCQ: None	<p>Description: A drive in the storage subsystem is uncertified.</p> <p>Action: Start the Recovery Guru to access the Uncertified Drive recovery procedure.</p>
Event 2602 - Automatic controller firmware synchronization failed	02/04/81	<p>Description: The versions of firmware on the redundant controllers are not the same because the automatic controller firmware synchronization failed. Controllers with an incompatible version of the firmware might cause unexpected results.</p> <p>Action: Try the firmware download again. If the problem remains, contact IBM support.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2801 - Storage subsystem running on uninterruptible power supply battery	6/3F/C8	<p>Description: The uninterruptible power supply has indicated that ac power is no longer present and the uninterruptible power supply has switched to standby power. There is no immediate cause for concern, but you must save your data frequently, in case the battery is suddenly depleted.</p> <p>Action: Start the Recovery Guru and click the Lost AC Power recovery procedure. Follow the instructions to correct the failure.</p>
Event 2803 - Uninterruptible power supply battery - two minutes to failure	6/3F/C9	<p>Description: The uninterruptible power supply has indicated that its standby power supply is nearing depletion.</p> <p>Action: Take actions to stop I/O activity to the controller. Normally, the controller changes from a write-back caching mode to a write-through mode.</p>
Event 2804 - Uninterruptible power supply battery failed	None	<p>Description: The uninterruptible power supply battery has failed.</p> <p>Action: Contact IBM support for the instructions for recovering from this failure.</p>
Event 2807 - Environmental service module failed	None	<p>Description: An ESM has failed.</p> <p>Action: Start the Recovery Guru and click the Failed Environmental Service Module CRU recovery procedure. Follow the instructions to correct the failure.</p>
Event 2808 - storage expansion enclosure ID not unique	6/98/01	<p>Description: The controller has determined that there are multiple storage expansion enclosures with the same ID selected. Verify that each storage expansion enclosure has a unique ID setting.</p> <p>Action: Start the Recovery Guru and click the Enclosure ID Conflict recovery procedure. Follow the instructions to correct the failure.</p>
Event 280A - Controller enclosure component missing	6/3F/C7	<p>Description: A component other than a controller is missing in the controller enclosure (for example, a fan, power supply, or battery). The FRU codes indicate the faulty component.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 280B - Controller enclosure component failed	6/3F/C7	<p>Description: A component other than a controller has failed in the controller enclosure (for example, a fan, power supply, battery), or an over-temperature condition has occurred. The FRU codes indicate the faulty component.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 280D - Drive storage expansion enclosures component failed	6/3F/C7	<p>Description: A component other than a drive has failed in the storage expansion enclosure (for example, a fan, power supply, or battery), or an over-temperature condition has occurred. The FRU codes indicate the faulty component.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 280E - Standby power supply not fully charged	6/3F/CA	<p>Description: The uninterruptible power supply has indicated that its standby power supply is not at full capacity.</p> <p>Action: Check the uninterruptible power supply to make sure that the standby power source (battery) is in working condition.</p>
Event 280F - Environmental service module - loss of communication	6/E0/20	<p>Description: Communication has been lost to one of the dual ESM CRUs in a storage expansion enclosure. The storage expansion enclosure has only one available I/O path.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2813 - Minihub CRU failed	6/3F/C7	<p>Description: Communication with the minihub CRU has been lost. This might be the result of a minihub CRU failure, a controller failure, or a failure in an internal backplane communications board. If there is only one minihub failure, the storage subsystem is still operational, but a second minihub failure might cause the failure of the affected enclosures.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2815 - GBIC failed	None	<p>Description: A gigabit interface converter (GBIC) on either the controller enclosure or the storage expansion enclosure has failed. If there is only one GBIC failure, the storage subsystem is still operational, but a second GBIC failure might cause the failure of the affected enclosures.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2816 - storage expansion enclosure ID conflict - duplicate IDs across storage expansion enclosures	6/98/01	<p>Description: Two or more storage expansion enclosures are using the same enclosure identification number.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2818 - storage expansion enclosure ID mismatch - duplicate IDs in the same storage expansion enclosure	6/98/02	<p>Description: A storage expansion enclosure in the storage subsystem contains ESMs with different enclosure identification numbers.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 281B - Nominal temperature exceeded	6/98/03	<p>Description: The nominal temperature of the enclosure has been exceeded. Either a fan has failed or the temperature of the room is too high. If the temperature of the enclosure continues to rise, the affected enclosure might automatically shut down. Fix the problem immediately, before it becomes more serious. The automatic shutdown conditions depend on the model of the enclosure.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 281C- Maximum temperature exceeded	6/3F/C6	<p>Description: The maximum temperature of the enclosure has been exceeded. Either a fan has failed or the temperature of the room is too high. This condition is critical and might cause the enclosure to shut down if you do not fix the problem immediately. The automatic shutdown conditions depend on the model of the enclosure.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 281D - Temperature sensor removed	6/98/03	<p>Description: A fan CRU that contains a temperature sensor has been removed from the storage subsystem.</p> <p>Action: Replace the CRU as soon as possible. Start the Recovery Guru, click the Failed or Removed Fan CRU recovery procedure, and follow the instructions to correct the failure.</p>
Event 281E - Environmental service module firmware mismatch	6/98/03	<p>Description: A storage expansion enclosure in the storage subsystem contains ESMs with different versions of firmware. ESMs in the same storage expansion enclosure must have the same version firmware. If you do not have a replacement service monitor, call IBM support to perform the firmware download.</p> <p>Action: Start the Recovery Guru and click the Environmental Service Module Firmware Version Mismatch recovery procedure. Follow the instructions to correct the failure.</p>
Event 2821 - Incompatible Minihub	None	<p>Description: An incompatible minihub blade has been detected in the controller enclosure.</p> <p>Action: Start the Recovery Guru and click the Incompatible minihub blade recovery procedure. Follow the instructions to correct the failure.</p>
Event 2823 - Drive bypassed	None	<p>Description: The ESM has reported that the drive has been bypassed to maintain the integrity of the Fibre Channel loop.</p> <p>Action: Start the Recovery Guru to access the By-Passed Drive recovery procedure and follow the instructions to recover from this failure.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2827 - Controller was inadvertently replaced with an ESM	None	<p>Description: A controller blade was inadvertently replaced with an ESM blade.</p> <p>Action: Replace the ESM blade with the controller blade as soon as possible.</p>
Event 2828 - Unsupported storage expansion enclosure selected	None	<p>Description: Your storage subsystem contains one or more unsupported storage expansion enclosures. If all of your storage expansion enclosures are being detected as being unsupported, you might have a problem with an NVSRAM configuration file or you might have the wrong version of firmware. This error condition will cause the drives in the unsupported storage expansion enclosures to be locked out, which can cause the defined arrays or logical drives to fail.</p> <p>Action: If there are array or logical drive failures, call IBM support for the recovery procedure. Otherwise, Start the Recovery Guru to access the Unsupported Drive Enclosure recovery procedure and follow the instructions to recover from this failure.</p>
Event 2829 - Controller redundancy lost	6/E0/20	<p>Description: Communication has been lost between the two controllers through one of the drive loops (channels).</p> <p>Action: Start the Recovery Guru and determine whether other loss of redundancy problems are being reported. If other problems are being reported, fix those first. If redundancy problems continue to be reported, contact IBM support.</p>
Event 282B - storage expansion enclosure path redundancy lost	6/E0/20	<p>Description: A storage expansion enclosure with redundant drive loops (channels) has lost communication through one of its loops. The storage expansion enclosure has only one loop that is available for I/O. Correct this failure as soon as possible. Although the storage subsystem is still operational, a level of path redundancy has been lost. If the remaining drive loop fails, all I/O to that storage expansion enclosure fails.</p> <p>Action: Start the Recovery Guru and click the Drive - Loss of Path Redundancy recovery procedure. Follow the instructions to correct the failure.</p>
Event 282D - Drive path redundancy lost	6/E0/20	<p>Description: A communication path with a drive has been lost. Correct this failure as soon as possible. The drive is still operational, but a level of path redundancy has been lost. If the other port on the drive or any other component fails on the working channel, the drive fails.</p> <p>Action: Start the Recovery Guru and click the Drive - Loss of Path Redundancy recovery procedure. Follow the instructions to correct the failure.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 282F - Incompatible version of ESM firmware detected	None	<p>Description: A storage expansion enclosure in the storage subsystem contains ESM blades with different firmware versions. This error might also be reported if a storage expansion enclosure in the storage subsystem contains ESM blades with different hardware.</p> <p>Action: Start the Recovery Guru to access the ESM blade Firmware Version Mismatch recovery procedure and follow the instructions to recover from this failure.</p>
Event 2830 - Mixed drive types not supported	None	<p>Description: The storage subsystem currently contains drives of different drive technologies, such as Fibre Channel (FC) and Serial ATA (SATA). Mixing different drive technologies is not supported on this storage subsystem.</p> <p>Action: Select the Recovery Guru to access the Mixed Drive Types Not Supported recovery procedure and follow the instructions to recover from this failure.</p>
Event 2835 - Drive storage expansion enclosures not cabled together	ASC/ASCQ: None	<p>Description: There are drive storage expansion enclosures in the storage subsystem that are not cabled correctly; they have ESM blades that must be cabled together sequentially.</p> <p>Action: Start the Recovery Guru to access the Drive Enclosures Not Cabled Together recovery procedure and follow the instructions to recover from this failure.</p>
Event 3019 - Logical drive ownership changed due to failover	None	<p>Description: The multipath driver software has changed ownership of the logical drives to the other controller because it could not access the logical drives on the particular path.</p> <p>Action: Start the Recovery Guru and click the Logical Drive Not on Preferred Path recovery procedure. Follow the instructions to correct the failure.</p>
Event 4011 - Logical drive not on preferred path	None	<p>Description: The controller that is listed in the Recovery Guru area cannot be accessed. Any logical drives for which this controller is assigned as their preferred path will be moved to the non-preferred path (alternative controller).</p> <p>Action: Start the Recovery Guru and click the Logical Drive Not on Preferred Path recovery procedure. Follow the instructions to correct the failure.</p>
Event 5005 - Place controller offline	None	<p>Description: The controller is placed offline. This might be caused by the controller failing a diagnostic test. (The diagnostics are initiated internally by the controller or by the Controller → Run Diagnostics menu option.) Or the controller is manually placed Offline using the Controller → Place Offline menu option.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 502F - Missing logical drive deleted	None	<p>Description: The storage subsystem has detected that the drives that are associated with a logical drive are no longer accessible. This can be the result of removing all the drives that are associated with an array or a loss of power to one or more storage expansion enclosures.</p> <p>Action: Start the Recovery Guru and click the Missing Logical Drive recovery procedure. Follow the instructions to correct the failure.</p>
Event 5038 - Controller in lockout mode	None	<p>Description: Both controllers have been placed in lockout mode for 10 minutes because password authentication failures have exceeded 10 attempts within a 10-minute period. During the lockout period, both controllers will deny all authentication requests. When the 10-minute lockout expires, the controller resets the total authentication failure counter and unlocks itself.</p> <p>Action: Wait 10 minutes and try to enter the password again.</p>
Event 5040 - Place controller in service mode	None	<p>Description: The controller was manually placed in service mode for diagnostic or recovery reasons.</p> <p>Action: Start the Recovery Guru to access the Controller in Service Mode recovery procedure. Use this procedure to place the controller back online.</p>
Event 5405 - Gold Key - mismatched settings	ASC/ASCQ: None	<p>Description: Each controller in the controller pair has a different NVSRAM bit setting that determines whether the controller is subject to Gold Key restrictions.</p> <p>Action: This event might be generated if IBM storage subsystem controllers or drives are inadvertently swapped with non-IBM controllers or drives. This critical event does not apply to the IBM DS3000, DS4000, or DS5000 storage subsystem configuration. Contact IBM support for the recovery procedure.</p>
Event 5406 - Mixed drive types - mismatched settings	ASC/ASCQ: None	<p>Description: Each controller in the controller pair has a different setting for the NVSRAM bit that controls whether Mixed Drive Types is a premium feature.</p> <p>Action: Start the Recovery Guru to access the Mixed Drive Types - Mismatched Settings recovery procedure and follow the instructions to correct this controller condition.</p>
Event 5602 - This controller alternate failed - timeout waiting for results	None	<p>Description: This controller initiated diagnostics on the alternative controller but did not receive a reply that indicates that the diagnostics were completed. The alternative controller in this pair has been placed offline.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 560B - CtlrDiag task cannot obtain Mode Select lock	None	<p>Description: This controller is attempting to run diagnostics and could not secure the test area from other storage subsystem operations. The diagnostics were canceled.</p> <p>Action: Contact IBM support for the instructions for recovering from this failure.</p>
Event 560C - CtlrDiag task on controller alternate cannot obtain Mode	None	<p>Description: The alternative controller in this pair is attempting to run diagnostics and could not secure the test area from other storage subsystem operations. The diagnostics were canceled.</p> <p>Action: Contact IBM support for the instructions for recovering from this failure.</p>
Event 560D - Diagnostics read test failed on controller	None	<p>Description: While the controller was running diagnostics, it detected that the information that was received does not match the expected return for the test. This might indicate that I/O is not being completed or that there is a mismatch in the data that is being read. The controller is placed offline as a result of this failure.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>
Event 560E - This controller alternate failed diagnostics read test	None	<p>Description: While the alternative for this controller was running diagnostics, it detected that the information that was received does not match the expected return for the test. This might indicate that I/O is not being completed or that there is a mismatch in the data that is being read. The alternative controller in this pair is placed offline.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>
Event 560F - Diagnostics write test failed on controller	None	<p>Description: While the alternative for this controller was running diagnostics, it was unable to write data to the test area. This might indicate that I/O is not being completed or that there is a mismatch in the data that is being written. The controller is placed offline.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>
Event 5610 - This controller alternate failed diagnostics write test	None	<p>Description: While the alternative for this controller was running diagnostics, it was unable to write data to the test area. This might indicate that I/O is not being completed or that there is a mismatch in the data that is being written. The alternative controller in this pair is placed offline.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 5616 - Diagnostics rejected - configuration error on controller	None	<p>Description: The alternative for this controller is attempting to run diagnostics and could not create the test area that is necessary for the completion of the tests. The diagnostics were canceled.</p> <p>Action: Contact IBM support for the instructions for recovering from this failure.</p>
Event 5617 - Diagnostics rejected - configuration error on controller alternate	None	<p>Description: The alternative for this controller is attempting to run diagnostics and could not create the test area that is necessary for the completion of the tests. The diagnostics were canceled.</p> <p>Action: Contact IBM support for the instructions for recovering from this failure.</p>
Event 6101 - Internal configuration database full	None	<p>Description: Because of the amount of data that is required to store certain configuration data, the maximum number of logical drives has been underestimated. One or both of the following types of data might have caused the internal configuration database to become full:</p> <ul style="list-style-type: none"> • FlashCopy logical drive configuration data • Global/Metro remote mirror configuration data <p>Action: To recover from this event, you can delete one or more FlashCopy logical drives from your storage subsystem, or you can remove one or more remote mirror relationships.</p>
Event 6107 - The alternate for the controller is nonfunctional and is being held in reset	None	<p>Description: A controller in the storage subsystem has detected that its alternative controller is nonfunctional because of hardware problems and must be replaced.</p> <p>Action: Start the Recovery Guru to access the Offline Controller recovery procedure and follow the instructions to recover from this failure.</p>
Event 6200 - FlashCopy repository logical drive threshold exceeded	None	<p>Description: The FlashCopy repository logical drive capacity has exceeded a warning threshold level. If the capacity of the FlashCopy repository logical drive becomes full, its associated FlashCopy logical drive can fail. This is the last warning that you receive before the FlashCopy repository logical drive becomes full.</p> <p>Action: Start the Recovery Guru and click the FlashCopy Repository Logical Drive Threshold Exceeded recovery procedure. Follow the instructions to correct this failure.</p>
Event 6201 - FlashCopy repository logical drive full	None	<p>Description: All of the available capacity on the FlashCopy repository logical drive has been used. The failure policy of the FlashCopy repository logical drive determines what happens when the FlashCopy repository logical drive becomes full. The failure policy can be set to either fail the FlashCopy logical drive (default setting) or fail incoming I/Os to the base logical drive.</p> <p>Action: Start the Recovery Guru and click the FlashCopy Repository Logical Drive Capacity - Full recovery procedure. Follow the instructions to correct this failure.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 6202 - Failed FlashCopy logical drive	None	<p>Description: Either the FlashCopy repository logical drive that is associated with the FlashCopy logical drive is full or its associated base or FlashCopy repository logical drives have failed because of one or more drive failures on their arrays.</p> <p>Action: Start the Recovery Guru and click the Failed FlashCopy Logical Drive recovery procedure. Follow the instructions to correct this failure.</p>
Event 6400 - Dual primary logical drive	None	<p>Description: Both logical drives have been promoted to primary logical drives after a forced role reversal. This event might be reported when the controller resets or when a cable from an array to a Fibre Channel switch is reinserted after it was removed and the other logical drive was promoted to a primary logical drive.</p> <p>Action: Start the Recovery Guru and click the Dual Primary Logical Drive Conflict recovery procedure. Follow the instructions to correct this failure.</p>
Event 6401 - Dual secondary logical drive	None	<p>Description: Both logical drives in the remote mirror have been demoted to secondary logical drives after a forced role reversal. This event might could be reported when the controller resets or when a cable from an array to a Fibre Channel switch is reinserted after it was removed and the other logical drive was promoted to a secondary logical drive.</p> <p>Action: Start the Recovery Guru and click the Dual Secondary Logical Drive Conflict recovery procedure. Follow the instructions to correct this failure.</p>
Event 6402 - Mirror data unsynchronized	Not recorded with event	<p>Description: This might occur because of I/O errors, but other events can be associated with it. A Needs Attention icon is displayed on both the primary and secondary storage subsystems of the remote mirror.</p> <p>Action: Start the Recovery Guru and click the Mirror Data Unsynchronized recovery procedure. Follow the instructions to correct this failure.</p>
Event 6503 - Remote logical drive link down	None	<p>Description: This event is triggered when a cable between one array and its peer has been disconnected, the Fibre Channel switch has failed, or the peer array has reset. This error might cause the Mirror Data Unsynchronized, event 6402. The affected remote logical drive displays an Unresponsive icon, and this state is selected in the tooltip when you pass your cursor over the logical drive.</p> <p>Action: Start the Recovery Guru and click the Mirror Communication Error - Unable to Contact Logical Drive recovery procedure. Follow the instructions to correct this failure.</p>

Table 38. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 6505 - WWN change failed	None	<p>Description: Mirroring causes a WWN change to be communicated between arrays. Failure of a WWN change is caused by non-I/O communication errors between one array, on which the WWN has changed, and a peer array. (The array WWN is the unique name that is used to locate an array on a fiber network. When both controllers in an array are replaced, the array WWN changes). The affected remote logical drive displays an Unresponsive icon and this state is selected in the tooltip when you pass your cursor over the logical drive.</p> <p>Action: Start the Recovery Guru and click the Unable to Update Remote Mirror recovery procedure. Follow the instructions to correct this failure. The only solution to this problem is to delete the remote mirror and then to establish another one.</p>
Event 6600 - Logical drive copy operation failed	None	<p>Description: A logical drive copy operation with a status of In Progress has failed. This failure can be caused by a read error on the source logical drive, a write error on the target logical drive, or a failure that occurred on the storage subsystem that affects the source logical drive or target logical drive.</p> <p>Action: Start the Recovery Guru and click the Logical Drive Copy Operation Failed recovery procedure. Follow the instructions to correct this failure.</p>
Event 6700 - Unreadable sector(s) detected - data loss occurred	None	<p>Description: Unreadable sectors have been detected on one or more logical drives, and data loss has occurred.</p> <p>Action: Start the Recovery Guru to access the Unreadable Sectors Detected recovery procedure and follow the instructions for recovering from this failure.</p>
Event 6703 - Overflow in unreadable sector database	None	<p>Description: The Unreadable Sectors log has been filled to its maximum capacity.</p> <p>Action: Select the Recovery Guru to access the Unreadable Sectors Log Full recovery procedure and follow the instructions for recovering from this failure.</p>

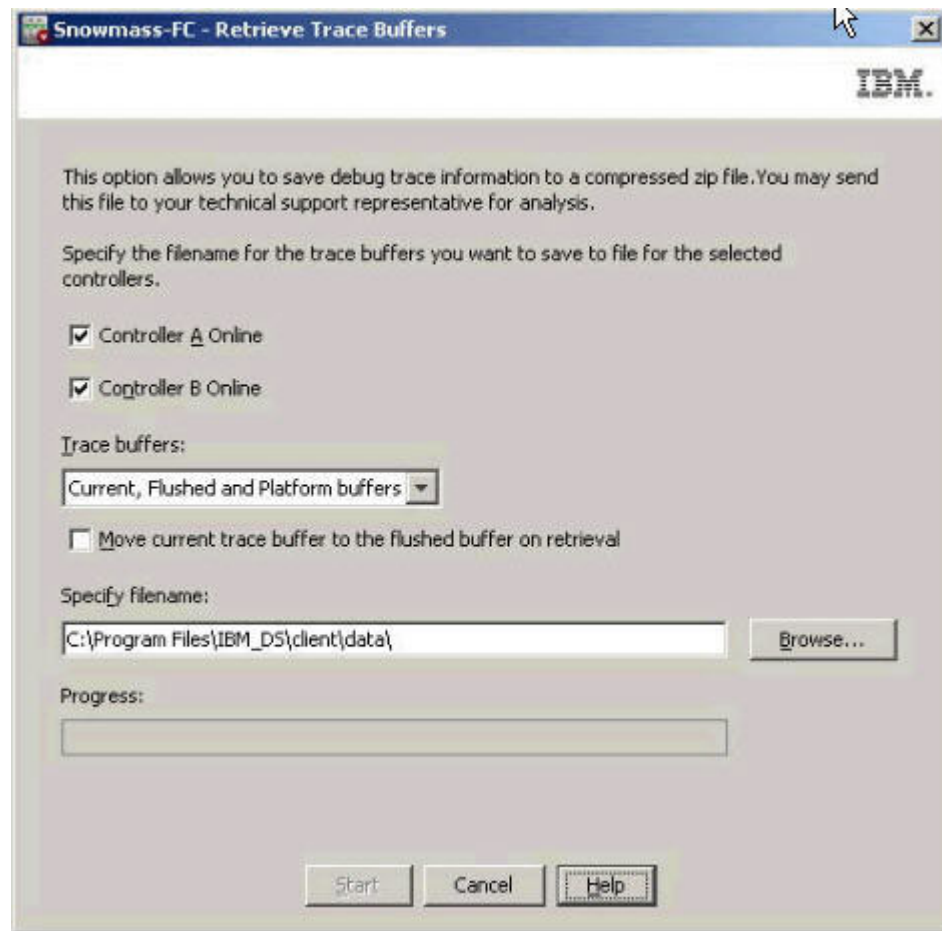
Retrieve trace buffers

Binary trace buffers provide data over a longer period of time than previous dqprint text in stateCaptureData.txt.

Advanced troubleshooting and support data are collected in binary format and must be parsed by IBM Support. They are a part of the Collect All Support Data (CASD) support bundle (traceBuffers.zip).

Usage in the CLI:

```
start controller [both] trace dataType=all forceFlush=FALSE file="C:\TBTest2.zip";
```



Configuration database validation

Version 10.77 adds a new Configuration Database Diagnostic feature.

To perform the validation manually, select a controller on the Physical Tab, then select **Advanced > Troubleshooting > Run Diagnostics > Configuration Database**.

An automatic database check will be performed before the controller firmware download is started through the EMW, SubSystem Managment Window, or CLI.

If validation fails, a zip file is created containing a text file with a description of the error and a DQ file for troubleshooting.


Validation failure MEL event:

Event Name: Raw data retrieve structure inconsistency detected

Type: 3408

Component Type: Controller

New CLI command: **start storageSubsystem configDbDiagnostic;**



Database save/restore

Storage Monitor Service automatically saves the configuration DB from a subsystem, and an existing configuration DB can be restored as well.

Save

Storage Monitor Service automatically saves the configuration DB from a subsystem and saves the file in "...client\data\monitor\dbcapture" if there is a DB change AND 125 minutes have passed since previous capture.

When a subsystem is added to a newly installed HSW, the first DB is captured.

All captured DB files are zipped and named as follows:

RetrievedRecords_SSID_Date_Time.dbm.

Example:

RetrievedRecords_60080e500017b8de000000004be47b12_2010_08_20_14_48_27.dbm

CLI can be used to save a DB manually by using the command **save storageSubsystem dbmDatabase file="C:\path\filename.zip"**

Restore

An existing configuration DB can be restored to recover systems that have lost their configuration or their configuration was removed to recover from failure.

It restores portions of the database containing:

- Lun and array configuration
- Lun WWNs
- Controller WWNs
- premium features
- mappings

It excludes:

- MEL

- UTM
- cache

Duration: up to 45 minutes

The user must have a Validator String to restore the configuration DB. To obtain the validator, send the config DB zip file and the system profile to IBM support. IBM support generates the validator string based on the information that you provide.

Load configuration database

To load the configuration database, do the following:

1. Open the *.key file sent via e-mail in a text editor.
2. Record the string. Example: 8bbaadfa7375cb4dfcc81c15bde30ad32d814c09
3. Stop the IO to the subsystem
4. Set one of the controllers offline via the GUI
5. The restore is done using the CLI command **load storageSubsystem dbmDatabase file="C:\path\filename.zip" validator="8bbaadfa7375cb4dfcc81c15bde30ad32d814c09";**

DS Diagnostic Data Capture (DDC)

The DDC function was implemented to assist IBM support in collecting data for troubleshooting unusual controller-firmware events.

Note: This function is not implemented with controller firmware code versions that are earlier than the 06.12.27.xx level.

Under rare circumstances, an internal controller error can force a routine to perform the Diagnostic Data Capture (DDC) function. When this occurs, a red stop sign next to the name of the storage subsystem that has the error (is in a non-optimal state) is displayed in the Enterprise Management window. After you open the Subsystem Management window for that storage subsystem, you can click the Recovery Guru. The Recovery Guru shows what the issue is, as does the MEL (Storage Manager Major Events Log). See “DDC MEL events” on page 265 for more information about the MEL.

When the DDC function is implemented, the storage subsystem status changes from Optimal to Needs Attention due to DDC. This occurs under the following conditions:

- The controllers in the storage subsystem detect unusual events such as Master Abort (because of a bad address that is accessed by the Fibre Channel chip, resulting in a PCI bus error).
- The controller is not able to process host I/O requests for an extended period of time (several minutes).
- The destination device number registry is corrupted.
- An EDC (error detection code) error is returned by the disk drives.
- A quiescence failure occurred in the logical drive that is owned by the alternative controller.
- The records that are related to Storage Partition Management are corrupted.

When the **Needs Attention due to DDC** flag is set, it is persistent across the power-cycle and controller reboot, provided that the controller cache batteries are sufficiently charged. In addition, data that reflects the state of the storage subsystem controllers at the moment that the unusual event occurred is collected and saved until you retrieve it. To clear the **Needs Attention due to DDC** flag and to retrieve the saved diagnostic data, see “Recovery steps.”

Because the current DDC function implementation holds the DDC data for only one unusual event at a time until the DDC data is saved, the **SMcli** commands must be performed as soon as possible whenever the Needs Attention due to DDC error occurs, so that the controllers can be ready for capturing data for any other unusual events. Until the diagnostic data is saved and the **Needs Attention due to DDC** flag is cleared, any occurrences of other unusual events do not trigger the controller to capture diagnostic data for those events. An unusual event invokes a DDC trigger if a previous DDC trigger is at least 48 hours old you have successfully retrieved the previous DDC information. In addition, DDC information is available only if a controller is online. A controller that is in service or lock-down mode does not trigger a DDC event. After you collect the DDC data, contact IBM support to report the problem and get assistance with troubleshooting the condition.

Recovery steps

To perform the DDC recovery process, complete the following steps:

1. Open either the Script Editor from the Enterprise Management window or the command-line interface (CLI).

Note: See the online help in the Enterprise Management window for more information about the syntax of these commands.

2. Follow the instructions in the following table, depending on whether you want to save the diagnostic data.

Table 39. Recovery Step 2

If...	Then...
You want to save the diagnostic data	Go to step 3.
You do not want to save the diagnostic data	Go to step 5.

3. Type


```
save storageSubsystem diagnosticData file="filename ";
```

 where *filename* is the location and name of the file that will be saved. The file is initialized as a .zip file.

Note: The *esm* parameter of the command syntax is not supported.

4. Follow the instructions in the following table to work with the diagnostic data.

Table 40. Recovery Step 4

If...	Then...
No error was returned	Go to step 6.

Table 40. Recovery Step 4 (continued)

If...	Then...	
An error was returned	If...	Then...
	The error message indicates that there was a problem saving the data.	Wait 2 minutes, and restart step 3.
	The error message indicates that there was a problem resetting the data.	Wait 2 minutes, and go to step 5.

5. Type

```
reset storageSubsystem diagnosticData;
```

Table 41. Recovery Step 5

If...	Then...
No error was returned	Go to step 6.
An error was returned	Wait 2 minutes and then run the command again. The controllers might need additional time to update the status. Note: Another error might occur if the diagnostic data status has already been reset. Go to step 6.

- Click **Recheck** to run the Recovery Guru again. The failure is no longer displayed in the **Summary** area.

After this process has been completed, the DDC message is removed automatically, and a recheck of the Recovery Guru shows no entries for DDC capture. If for some reason the data has not been removed, the Recovery Guru provides an example of how to clear the DDC information without saving the data. To complete the preceding procedure in the script editor, type

```
reset storageSubsystem diagnosticData;
```

DDC MEL events

When the Diagnostic Data Capture action is triggered by an unusual event, one or more of the following events are posted in the storage subsystem event logs, depending on the user actions.

Table 42. DDC MEL events

Event number	Description	Priority	Explanation
0x6900	Diagnostic Data is available.	Critical	This is logged when an unusual controller event triggers the DDC function to store Diagnostic Data.
0x6901	Diagnostic Data retrieval operation started.	Informational	This is logged when the user runs the SMcli command to retrieve and save the Diagnostic Data, as described in step 3 in "Recovery steps" on page 264.
0x6902	Diagnostic Data retrieval operation completed.	Informational	This is logged when the Diagnostic Data retrieval and save completes.

Table 42. DDC MEL events (continued)

Event number	Description	Priority	Explanation
0x6903	Diagnostic Data Needs Attention status/flag cleared.	Informational	This is logged when the user resets the <i>Needs Attention due to DDC</i> flag with the SMcli command, or when the Diagnostic Data retrieval and save completes successfully when initiated by the user with the save storageSubsystem diagnosticData SMcli command.

Resolving disk array errors on AIX

This section describes the disk array errors that might be reported in the AIX error log. You can view the AIX error log by using the **errpt -a** command. You can also check the Storage Manager Major Event log (MEL) to find out whether there is any correlation between the host, SAN, and storage subsystem.

You might have to validate your configuration or replace defective hardware to correct the situation.

Note: For more troubleshooting information, see the *Installation, User's, and Maintenance Guide* that came with your storage subsystem.

Table 43. Disk array errors

Error number	Error name	Error type	Error description
1	FCP_ARRAY_ERR1	ARRAY OPERATION ERROR	A permanent hardware error involving the disk array media occurred.
2	FCP_ARRAY_ERR2	ARRAY OPERATION ERROR	A permanent hardware error occurred.
3	FCP_ARRAY_ERR3	ARRAY OPERATION ERROR	A permanent error was detected by the array adapter.
4	FCP_ARRAY_ERR4	ARRAY OPERATION ERROR	A temporary error occurred within the array, communications, or adapter.
5	FCP_ARRAY_ERR5	UNDETERMINED ERROR	An undetermined error occurred.
6	FCP_ARRAY_ERR6	SUBSYSTEM COMPONENT FAILURE	A degradation condition other than a disk drive occurred.
7	FCP_ARRAY_ERR7	CONTROLLER HEALTH CHECK FAILURE	A health check on the passive controller failed.
8	FCP_ARRAY_ERR8	ARRAY CONTROLLER SWITCH	One array controller became unavailable, so I/O moved to the other controller.
9	FCP_ARRAY_ERR9	ARRAY CONTROLLER SWITCH FAILURE	An array controller switch failed.

Table 43. Disk array errors (continued)

Error number	Error name	Error type	Error description
10	FCP_ARRAY_ERR10	ARRAY CONFIGURATION CHANGED	A logical unit was moved from one controller to the other (most likely by the action of an alternative host).
11	FCP_ARRAY_ERR11	IMPROPER DRIVE TYPE FOR DUAL ACTIVE MODE	This error is not possible on the 2102 array and exists for historical purposes only. FCP_ARRAY_ERR11 might be reused for a different error in the future.
12	FCP_ARRAY_ERR12	POLLED AEN FAILURE	An automatic error notification failed.
13	FCP_ARRAY_ERR13	ARRAY INTER-CONTROLLER COMMUNICATION FAILURE	The controllers are unable to communicate with each other. This error might occur if one of the controllers is rebooted while the error log is being generated. However, it might also indicate a problem with the Fibre Channel connections.
14	FCP_ARRAY_ERR14	ARRAY DRIVE FAILURE	A serious or unrecoverable error was detected on a physical disk within the storage subsystem. A system engineer might be able to obtain the exact cause from an analysis of the sense data.
15	FCP_ARRAY_ERR15	CACHE BATTERY LOW/DATA LOSS POSSIBLE	If a controller card is replaced, the cache batteries might be drained. It can take two days for the cache batteries to be recharged. During this time, errors are logged in the error log. Do not replace the controller.
16	FCP_ARRAY_ERR16	CACHE BATTERY CHARGE BELOW 87.5%	If a controller card is replaced, the cache batteries might be drained. It can take two days for the cache batteries to be recharged. During this time, errors are logged in the error log. Do not replace the controller.
17	FCP_ARRAY_ERR17	WORLDWIDE NAME CHANGED	A controller has changed worldwide names. This error might be caused if you replace the controller without placing it in the reset state first, or if you change the cabling and a different controller with the same SCSI ID is on the loop.
18	FCP_ARRAY_ERR18	RESERVATION CONFLICT	An operation failed because the disk array logical drive (LUN) is reserved by another host.

Table 43. Disk array errors (continued)

Error number	Error name	Error type	Error description
19	FCP_ARRAY_ERR19	SNAPSHOT VOLUME REPOSITORY FULL	The repository capacity limit was reached. To resolve this error, increase the repository capacity.
20	FCP_ARRAY_ERR20	SNAPSHOT OPERATION STOPPED BY ADMIN	The FlashCopy (snapshot) operation was disabled or stopped. To resolve this error, re-create the FlashCopy.
21	FCP_ARRAY_ERR21	SNAPSHOT REPOSITORY METADATA ERROR	There was a problem with the metadata of the FlashCopy (snapshot) repository during the FlashCopy operation. To resolve this error, re-create the FlashCopy.
22	FCP_ARRAY_ERR22	REMOTE VOL MIRRORING: ILLEGAL I/O ORIGIN	The primary logical drive received I/O from a remote array, or the secondary logical drive received I/O from a source other than the primary logical drive. To resolve this error, try the operation again.
23	FCP_ARRAY_ERR23	SNAPSHOT OPERATION NOT ALLOWED	The repository capacity limit was reached, so the FlashCopy (snapshot) operation failed. To resolve this error, delete or re-create the FlashCopy.
24	FCP_ARRAY_ERR24	SNAPSHOT VOLUME REPOSITORY FULL	The repository-capacity limit was reached. To resolve this error, delete or re-create the FlashCopy (snapshot).
25	FCP_ARRAY_ERR25	CACHED DATA WILL BE LOST IF CONTROLLER FAILS	<p>This message is a warning that a disk array logical drive (LUN) is running with write cache enabled and cache mirroring disabled. The warning is displayed when the LUN is opened, and it is displayed again every 24 hours until cache mirroring is enabled again.</p> <p>If a controller fails, or if power to the controller is turned off while the LUN is running in this mode, data that is in the write cache (but not written to the physical disk media) might be lost. This can cause corrupted files, file systems, or databases.</p>

Table 43. Disk array errors (continued)

Error number	Error name	Error type	Error description
26	FCP_ARRAY_ERR26	LOGICAL VOLUME IS WRITE PROTECTED	<p>The status of the logical drive is read-only. The probable reason is that it is a secondary logical drive of a FlashCopy, VolumeCopy, or remote mirror pair. Determine which relationship applies to the logical drive.</p> <ul style="list-style-type: none"> • For FlashCopy, a status of read-only on the secondary logical drive usually indicates that the repository is full. • For VolumeCopy, both the primary and secondary logical drives are read-only during the copy operation. The secondary logical drive is read-only when the copy operation is stopped and the copy pair has not been deleted. • For remote mirroring, the secondary logical drive is always read-only, provided that the mirror is active.
27	FCP_ARRAY_ERR27	SINGLE CONTROLLER RESTARTED	<p>The storage subsystem is operating as a single controller, and an error was repaired. The error might be caused by a communication or hardware problem, or it might occur because a LUN was moved to a controller that does not have a path to the current host.</p> <p>If this is a dual-controller storage subsystem, find the reason that the storage subsystem is operating in single-controller mode, and resolve the problem. Possible reasons include the following causes:</p> <ul style="list-style-type: none"> • An HBA, switch port, switch, storage subsystem port or storage subsystem controller was unavailable during the last system restart or the last time the cfgmgr command was run. • You removed a path (dac) as part of a Fibre Channel adapter hot-swap operation.

Table 43. Disk array errors (continued)

Error number	Error name	Error type	Error description
28	FCP_ARRAY_ERR28	SINGLE CONTROLLER RESTART FAILURE	<p>The storage subsystem is operating as a single controller, and the error has not been repaired. There is a problem with the path between this host and the storage subsystem or with the storage subsystem itself. The host has attempted to communicate with the storage subsystem, and that communication has failed.</p> <p>If the number of retries that is specified in the ODM attribute <code>switch_retries</code> is reached, the I/O is failed back to the user.</p> <p>Repair the error. Then, if this is a dual-controller storage subsystem, find the reason that the storage subsystem is operating in single-controller mode, and resolve that problem. Possible reasons include the following causes:</p> <ul style="list-style-type: none"> • An HBA, switch port, switch, storage subsystem port or storage subsystem controller was unavailable during the last system restart or the last time the cfgmgr command was run. • You removed a path (dac) as part of a Fibre Channel adapter hot-swap operation.

A new errorlog DISK_ERR7 has been created to notify that a path has been designated as failed because of a predetermined number of IO errors that occurred on the path. This is normally preceded with other error logs that represent the actual error that occurred on the path.

IBM DS Storage Manager - Password Reset

In case you have forgotten your password and are unable to log in to the IBM DS Storage Manager, you can press **Password Reset** on the controller panel and access the subsystem.

To know where the **Password Reset** button is located, see the *Installation, User's, and Maintenance Guide* for your storage subsystem.

Appendix A. Host bus adapter settings

This chapter covers the default settings for a variety of host bus adapters (HBAs) suitable for use with DS3000, DS4000, and DS5000 storage subsystems for Windows, Linux on Intel, VMware ESX, and NetWare operating systems. All other operating systems and platforms must use the default values. See the applicable product documentation for more information.

See the readme file that is included in the Fibre Channel host bus adapter BIOS or device driver package for any up-to-date changes to the settings.

An HBA is used to connect servers to Fibre Channel topologies. Its function is similar to that provided by network adapters to access LAN resources. The device driver for an HBA is typically responsible for providing support for a Fibre Channel topology, whether point-to-point, loop, or fabric.

Important: IBM supports the coexistence of DS3000, DS4000, and DS5000 storage subsystems in the same server, using the same HBAs, only if the application can accommodate the common set of device and multipath driver time-out settings. The common set of time-out settings consists of the longest settings that are published in this section for the DS3000, DS4000, and DS5000 storage subsystems. If this is not possible, and if the multipath driver allows for it, the use of mixed DS3000, DS4000, and DS5000 storage subsystems can still be achieved if the HBAs that you use for DS3000 subsystems are different than the ones that you use for DS4000 or DS5000 storage subsystems. Otherwise, the coexistence of DS3000, DS4000, and DS5000 storage subsystems in the same server is not possible.

See the latest IBM DS controller firmware and Storage Manager readme files for any restrictions on the coexistence of DS3000, DS4000, and DS5000 storage subsystems in the same server. For detailed HBA support information, see www.ibm.com/systems/support/storage/config/ssic.

Adjusting HBA settings

It is often necessary to adjust the settings of your HBA to match the capabilities of your device. This section describes how to access those settings to make the necessary adjustments.

Accessing HBA settings through Fast!UTIL

The Fast!UTIL feature provides access to host bus adapter settings. To access this feature, press Alt+Q or Ctrl+Q during BIOS initialization. It might take a few seconds for the Fast!UTIL menu to be displayed. If more than one adapter is installed, Fast!UTIL prompts you to select an adapter to configure. After you change adapter settings, Fast!UTIL restarts the server to load the new parameters. After you enter Fast!UTIL, the following selections are available on the **Fast!UTIL Options** menu:

- Configuration Settings
- Loopback Test
- Select Host Adapter

You can also access the host bus adapter settings through the **Configuration Settings** menu in Fast!UTIL; then, select **Adapter Settings** or **Advanced Adapter Settings**.

Note: Alternatively, you can also use the QLogic SANsurfer program to modify the **Host adapter settings** and **Advanced adapter settings** preferences from the Microsoft Windows operating-system environment. You must restart the servers for the changes to become effective.

Default host bus adapter settings

Access the host bus adapter settings through the **Configuration Settings** menu in Fast!UTIL and select **Adapter Settings**. The default host bus adapter settings for the FC2-133 HBA are as follows:

Host Adapter BIOS

When this setting is Disabled, the ROM BIOS on the FC2-133 HBA is disabled, and space becomes available in upper memory. This setting must be Enabled if you are booting from a Fibre Channel disk drive that is attached to the FC2-133 adapter. The default is Disabled.

Frame Size

This setting specifies the maximum frame length that is supported by the FC2-133 HBA. The default size is 2048, which provides maximum performance for F-Port (point-to-point) connections.

Loop Reset Delay

After resetting the loop the firmware refrains from initiating any loop activity for the number of seconds that is specified in this setting. The default is 5 seconds.

Adapter Hard Loop ID

This setting forces the adapter to attempt to use the ID that is specified in the Hard Loop ID setting. The default is Enabled.

Hard Loop ID

If the Adapter Hard Loop ID setting is Enabled, the adapter attempts to use the ID that is specified in this setting. The default ID is 125. Set this ID to a unique value from 0-125 if there is more than one adapter that is connected to an FC-AL loop and the Adapter Hard Loop ID setting is Enabled.

Spin Up Delay

When this bit is set, the BIOS waits up to 5 minutes to find the first drive. The default setting is Disabled.

Connection Options

This setting defines the type of connection (loop or point-to-point) or connection preference. The default is 2, which is loop preferred unless point-to-point.

Fibre Channel Tape Support

This setting enables FCP-2 recovery. The default is Enabled. Change this setting to Disabled if the HBA is not connected to a tape device.

Data Rate

This setting determines the data rate. When this setting is 0, the FC2-133 HBA runs at 1 Gbps. When this setting is 1, the FC2-133 HBA runs at 2 Gbps. When this setting is 2, Fast!UTIL determines what rate your system can accommodate and sets the rate accordingly. The default is 2 (auto-configure).

Advanced HBA settings

Access the following advanced host bus adapter settings through the **Configuration Settings** menu in Fast!UTIL and select **Advanced Adapter Settings**. The default settings for the FC2-133 HBA are as follows:

Execution Throttle

This setting specifies the maximum number of commands that execute on any one port. When a port execution throttle is reached, no new commands are executed until the current command is finished. The valid options for this setting are 1-256. The default is 255.

LUNs per Target

This setting specifies the number of LUNs per target. Multiple LUN support is typically for redundant array of independent disks (RAID) systems that use LUNs to map drives. The default is 0. For host operating systems other than Microsoft Windows, you might have to change this setting to a value other 0 to allow the host to see more than one logical drive from the storage subsystem.

Enable LIP Reset

This setting determines the type of loop initialization process (LIP) reset that is used when the operating system initiates a bus reset routine. When this setting is Yes, the driver initiates a global LIP reset to clear the target device reservations. When this setting is no, the driver initiates a global LIP reset with full login. The default is No.

Enable LIP Full Login

This setting instructs the ISP chip to log in, again, to all ports after any LIP. The default is Yes.

Enable Target Reset

This setting enables the drivers to issue a Target Reset command to all devices on the loop when a SCSI Bus Reset command is issued. The default is Yes.

Login Retry Count

This setting specifies the number of times that the software tries to log in to a device. The default is 30 retries.

Port Down Retry Count

This setting specifies the number of seconds that elapse before the software retries a command to a port returning port down status. The default is 30 seconds. For the Microsoft Windows servers in MSCS configuration, the Port Down Retry Count BIOS parameter must be changed from the default of 30 to 70.

Link Down Timeout

This setting specifies the number of seconds that the software waits for a link down to come up. The default is 60 seconds.

Extended Error Logging

This setting provides additional error and debug information to the operating system. When it is enabled, events are logged in the Windows NT Event Viewer. The default is Disabled.

RIO Operation Mode

This setting specifies the reduced interrupt operation (RIO) modes, if supported by the software driver. RIO modes allow posting multiple command completions in a single interrupt. The default is 0.

Interrupt Delay Timer

This setting contains the value (in 100-microsecond increments) used by a timer to set the wait time between accessing (DMA) a set of handles and generating an interrupt. The default is 0.

QLogic host bus adapter settings

Important: IBM supports the coexistence of DS3000, DS4000, and DS5000 storage subsystems in the same server, using the same HBAs, only if the application can accommodate the common set of device and multipath driver time-out settings. The common set of time-out settings consists of the longest settings that are published in this section for the DS3000, DS4000, and DS5000 storage subsystems. If this is not possible, and if the multipath driver allows for it, the use of mixed DS3000, DS4000, and DS5000 storage subsystems can still be achieved if the HBAs that you use for DS3000 subsystems are different than the ones that you use for DS4000 or DS5000 storage subsystems. Otherwise, the coexistence of DS3000, DS4000, and DS5000 storage subsystems in the same server is not possible.

See the latest IBM DS controller firmware and Storage Manager readme files for any restrictions on the coexistence of DS3000, DS4000, and DS5000 storage subsystems in the same server. For detailed HBA support information, see www.ibm.com/systems/support/storage/config/ssic.

Note: The BIOS settings in the Windows column are the default values that are set when the adapters are ordered from IBM as IBM FC-2 (QLA2310), FC2-133 (QLA2340) and single-port and dual-port 4 Gbps (QLx2460 and QLx2462) Fibre Channel host bus adapters. If the adapters are not from IBM, the default BIOS might not be the same as those defined in the Microsoft Windows column. There is one exception: the default setting for Fibre Channel tape support is enabled.

Table 44 shows the default settings for IBM Fibre Channel FC-2 and FC2-133 (QLogic adapter models QLA2310 and QLA2340) host bus adapter settings (for BIOS V1.35 and later) by operating system as well as the default registry settings for Microsoft Windows operating systems. DS3000, DS4000, or DS5000 products require BIOS V1.43 or later for these adapters. In addition, these settings are also the default BIOS settings for the newer DS3000, DS4000, or DS5000 4 Gbps single and dual-port host bus adapters (QLogic adapter models QLx2460 and QLx2462). The 4 Gbps host bus adapter BIOS version is 1.12 or later. See the applicable readme file for the latest updates to these values.

Table 44. QLogic model QLA234x, QLA24xx, QLE2462, QLE2460, QLE2560, QLE2562, QMI2572, QMI3572, QMI2582

Item	Default	VMware	Windows 2000	Windows 2003 and Windows 2008	Solaris	LINUX MPP	LINUX DMMP	NetWare
BIOS settings								
Host Adapter settings								
Host Adapter BIOS	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
Frame Size	2048	2048	2048	2048	2048	2048	2048	2048
Loop Reset Delay	5	5	8	8	8	8	8	8

Table 44. QLogic model QLA234x, QLA24xx, QLE2462, QLE2460, QLE2560, QLE2562, QMI2572, QMI3572, QMI2582 (continued)

Item	Default	VMware	Windows 2000	Windows 2003 and Windows 2008	Solaris	LINUX MPP	LINUX DMMP	NetWare
Adapter Hard Loop ID – (only for arbitrated loop topology).	Disabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Hard Loop ID (must be unique for each HBA) – (only for arbitrated loop topology).	0	125 ¹	125 ¹	125 ¹	125 ¹	125 ¹	125 ¹	125 ¹
Spin-up Delay	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
Connect Options	2	2	2	2	2	2	2	2
Fibre Channel Tape Support	Disabled	Disabled ³	Disabled ³	Disabled ³	Disabled ³	Disabled ³	Disabled ³	Disabled ³
Data Rate	2	2 (Auto)	2 (Auto)	2 (Auto)	2 (Auto)	2 (Auto)	2 (Auto)	2 (Auto)
Advance Adapter Settings								
Execution Throttle	16	256	256	256	256	256	256	256
LUNs per Target	8	0	0	0	0	0	0	32
Enable LIP Reset	No	No	No	No	No	No	No	No
Enable LIP Full Login	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Enable Target Reset	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Login Retry Count	8	30	30	30	30	30	30	30
Port Down Retry Count (5.30 controller firmware and earlier)	8	30	30	30	30	12	12	70
Port Down Retry Count	8	70	DS3K: 144 DS4K/5K: 70 ²	DS3K: 144 DS4K/5K: 70 ²	70	DS3K: 70 DS4K5K: 35	10	70
Link Down Timeout	30	60	DS3K:144 DS4K/5K: 60	DS3K:144 DS4K/5K: 60	60	DS3K:144 DS4K/5K: 60	NA	60
Extended Error Logging	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
RIO Operation Mode	0	0	0	0	0	0	0	0
Interrupt Delay Timer	0	0	0	0	0	0	0	0
IOCB Allocation	256	256	256	256	256	256	256	256
>4 GB Addressing	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
Drivers Load RISC Code	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Enable Database Updates	No	No	No	No	No	No	No	No
Disable Database Load	No	No	No	No	No	No	No	No
Fast Command Posting	Disabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled

Table 44. QLogic model QLA234x, QLA24xx, QLE2462, QLE2460, QLE2560, QLE2562, QMI2572, QMI3572, QMI2582 (continued)

Item	Default	VMware	Windows 2000	Windows 2003 and Windows 2008	Solaris	LINUX MPP	LINUX DMMP	NetWare
Extended Firmware Settings (1.34 and Earlier)								
Extended Control Block	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
RIO Operation Mode	0	0	0	0	0	0	0	0
Connection Options	2	2	2	2	2	2	2	2
Class 2 Service	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
ACK0	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
Fibre Channel Tape Support	Enabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
Fibre Channel Confirm	Enabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
Command Reference Number	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
Read Transfer Ready	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
Response Timer	0	0	0	0	0	0	0	0
Interrupt Delay Timer	0	0	0	0	0	0	0	0
Data Rate	2	2 (Auto)	2 (Auto)	2 (Auto)	2 (Auto)	2 (Auto)	2 (Auto)	2 (Auto)

REGISTRY SETTINGS⁵

(HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\QL2300\Parameters\Device)

LargeLuns	1	1	N/A	N/A	N/A	N/A	N/A	N/A
MaximumSGList	0x21	0xff	0xff	0xff	N/A	N/A	N/A	N/A

O/S REGISTRY SETTINGS⁵

(HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\QL2300\Parameters\Device) under DriverParameter variable.

Note:

1. Prior to QLogic driver versions 9.1.x.x, the variable name used was DriverParameters instead of DriverParameter.
2. DriverParameter is of type REG_SZ and the following parameters are added to the DriverParameters string. Do not create a separate key for each of the parameters.

UseSameNN	1	1	1	1	N/A	N/A	N/A	N/A
BusChange (SCSI Port Miniport 9.0.1.60 and earlier – does not apply to 9.1.1.11 and newer)	2	N/A	0	0	N/A	N/A	N/A	N/A
TimeoutValue 4 (REG_DWORD)	0x3C	N/A	DS3K: xA0 DS4K/5K: x78	DS3K: xA0 DS4K/5K: x78	N/A	N/A	N/A	N/A

REGISTRY SETTINGS⁵

(HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\<FAILOVER>\parameters: Where <FAILOVER>=Rdacdisk for MPP or RDAC installations or <FAILOVER>=mppdsm, ds4dsm, md3dsm, sx3dsm, csmdsm, or tpsdsm for MPIO installations. Mppdsm is for the generic version, your installation could be different.)

Table 44. QLogic model QLA234x, QLA24xx, QLE2462, QLE2460, QLE2560, QLE2562, QMI2572, QMI3572, QMI2582 (continued)

Item	Default	VMware	Windows 2000	Windows 2003 and Windows 2008	Solaris	LINUX MPP	LINUX DMMP	NetWare
SynchTimeOut (REG_DWORD)	0x78	N/A	DS3K: xA0 DS4K/5K: x78	DS3K: xA0 DS4K/5K: x78				
DisableLunRebalance (Only applies to cluster configurations. Firmware version 6.xx.xx.xx and later.)	0x00	N/A	0x03	0x03				

SuSE 7.3 specific modifications:

- Offset 0x11 in the Linux region (6) of the array controller NVSRAM must be changed from the default of 0x20 to 0x7f. The following command can be run from the script engine:
 - Set controller[a] HOSTNVSRAMByte[6,0x11]=0x7f;
 - Set controller[b] HOSTNVSRAMByte[6,0x11]=0x7f;
- The QLogic driver source must be modified to reflect the symbolic link used by SuSE.
 - vi makefile
 - find OSVER and change it from OSVER=linux-2.4 to OSVER=linux
 - Save and quit

Red Hat Linux Advanced Server 2.1 / SuSE Linux Enterprise Server 8.0 (6.x series failover driver [with no RDAC] only). Append the following to the HBA driver option string in the /etc/modules.conf file: ql2xretrycount=60 ql2xsuspendcount=40

If you are running the QLogic Inbox driver, the string options qla2xxx qlport_down_retry=144 (PB1-3) or options qla2xxx qlport_down_retry=70 (PB4-6) must be added in /etc/modprobe.conf (for RHEL) or /etc/modprobe.conf.local (for SLES). For all prior (RH3/4 SLES8/9) Linux versions (and out-of-box drivers), the string options qla2xxx qlport_down_retry=72 (PB1-3) or options qla2xxx qlport_down_retry=35 (PB4-6) must be added instead.

Note:

1. This setting must be changed to a unique AL-PA value if there is more than one Fibre Channel device in the FC-AL loop.
2. For larger configurations with heavy I/O loads or in a Microsoft cluster service (MSCS) environment, this value might be increased.
3. Change this setting to Enabled or Supported when the HBA is connected to a tape device only. Set it to Disabled when you connect to a DS3000, DS4000, or DS5000 storage subsystem.
4. In certain storage subsystem maximum configuration installations, you might have to set the TimeoutValue to 120 (decimal). Changing this value to a higher value might affect your application especially when it requires the disk I/O completion acknowledgement within a certain amount of time.
5. You can access registry settings by clicking **Start**, select **Run...**, type regedit into the **Open:** field, and then click **OK**.

Attention: Exercise caution when you change the Windows registry. If you change the wrong registry entry or make an incorrect entry for a setting, you can cause an error that prevents your server from booting or operating correctly.

Note: The BIOS settings under the Windows column are the default values that are set when the adapters are ordered from IBM as IBM Fibre Channel host bus adapters. If the adapters are not from IBM, the default BIOS might not be the same

as the ones that are defined in the Microsoft Windows column. There is one exception: the default setting for Fibre Channel tape support is enabled.

Table 45 shows the default settings for various IBM DS3000, DS4000, or DS5000 Fibre Channel host bus adapters (QLogic adapter QL220x) models (for BIOS V1.81) by operating system. See the applicable readme file for the latest updates to these values.

Table 45. QLogic model QL220x (for BIOS V1.81) host bus adapter settings by operating system

Item	Windows		Linux	NetWare
	NT	2000 / Server 2003		
BIOS settings				
Host Adapter settings				
Host Adapter BIOS	Disabled	Disabled	Disabled	Disabled
Frame Size	2048	2048	2048	2048
Loop Reset Delay	5	5	8	5
Adapter Hard Loop ID	Enabled	Enabled	Enabled	Enabled
Hard Loop ID (must be unique for each HBA)	125 ¹	125 ¹	125 ¹	125 ¹
Spin Up Delay	Disabled	Disabled	Disabled	Disabled
Advanced adapter settings				
Execution Throttle	256	256	256	256
>4 Gbyte Addressing	Disabled	Disabled	Disabled	Disabled
LUNs per Target	0	0	0	32
Enable LIP Reset	No	No	No	No
Enable LIP Full Login	Yes	Yes	Yes	Yes
Enable Target Reset	Yes	Yes	Yes	Yes
Login Retry Count	30	30	30	30
Port Down Retry Count	30	30	12	30 ²
IOCB Allocation	256	256	256	256
Extended Error Logging	Disabled	Disabled	Disabled	Disabled
Extended Firmware Settings				
Extended Control Block	Enabled	Enabled	Enabled	Enabled
RIO Operation Mode	0	0	0	0
Connection Options	3	3	3	3
Class 2 Service	Disabled	Disabled	Disabled	Disabled
ACK0	Disabled	Disabled	Disabled	Disabled
Fibre Channel Tape Support	Supported ³	Supported ³	Supported ³	Supported ³
Fibre Channel Confirm	Disabled	Disabled	Disabled	Disabled
Command Reference Number	Disabled	Disabled	Disabled	Disabled
Read Transfer Ready	Disabled	Disabled	Disabled	Disabled
Response Timer	0	0	0	0
Interrupt Delay Time	0	0	0	0
Registry settings ⁴ (HKEY_LOCAL_MACHINE → System → CurrentControlSet → Services → QL2200 → Parameters → Device)				

Table 45. QLogic model QL220x (for BIOS V1.81) host bus adapter settings by operating system (continued)

LargeLuns		1		
MaximumSGList	0x21	0x21		
Registry settings ⁴ (HKEY_LOCAL_MACHINE → System → CurrentControlSet → Services → Disk)				
TimeOutValue ⁴ (REG_DWORD)	0x3C	0x3C		
Registry settings ⁴ (HKEY_LOCAL_MACHINE → System → CurrentControlSet → Services → QL2200 → Parameters → Device) under the DriverParameter variable				
BusChange		0		

Note:

1. This setting must be changed to a unique AL-PA value if there is more than one Fibre Channel device in the FC-AL loop.
2. For larger configurations with heavy I/O loads, change this value to 70.
3. Change this setting to Enable or Supported when the HBA is connected to a tape device only. Set it to Disabled when you connect to DS3000, DS4000, or DS5000 Storage Subsystem.
4. To access registry settings, click **Start**, select **Run**, type regedit into the **Open** field, and then click **OK**.

Attention: Exercise caution when you change the Windows registry. If you change the wrong registry entry or make an incorrect entry for a setting, you can cause an error that prevents your server from booting or operating correctly.

JNI and QLogic host bus adapter settings

The following tables detail settings for the various host bus adapters (HBA) for Sun Solaris.

Note: JNI host bus adapters are supported only on Solaris 8 and 9. They are not supported on Solaris 10.

JNI HBA card settings

The JNI cards are not Plug and Play with autoconfiguration. Instead, you might have to change the settings or bindings.

Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562

JNI host bus adapters models FCE-1473, FCE-6460, FCX2-6562, and FCC2-6562 are supported with all currently supported levels of storage subsystem controller firmware.

Important: For each setting that is listed in Table 46, you must uncomment the line. This is true both for default settings and for settings that you must change.

Table 46. Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562

Original value	New value
FcLoopEnabled = 1	FcLoopEnabled = 0 (for non-loop; auto-topology) FcLoopEnabled = 1 (for loop)
FcFabricEnabled = 0	FcFabricEnabled = 0 (for non-fabric; auto-topology) FcFabricEnabled = 1 (for fabric)
FcEngHeartbeatInterval = 5	Same as original value (in seconds)

Table 46. Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562 (continued)

Original value	New value
FcLinkUpRecoveryTime = 1000	Same as original value (in milliseconds)
BusRetryDelay = 5000	Same as original value (in milliseconds)
TargetOfflineEnable = 1	TargetOfflineEnable = 0 (Disable) TargetOfflineEnable = 1 (Enable)
FailoverDelay = 30;	FailoverDelay = 60 (in seconds)
FailoverDelayFcTape = 300	Same as original value (seconds)
TimeoutResetEnable = 0	Same as original value
QfullRetryCount = 5	Same as original value
QfullRetryDelay = 5000	Same as original value (in milliseconds)
LunRecoveryInterval = 50	Same as original value (in milliseconds)
FcLinkSpeed = 3	Same as original value
JNICreationDelay = 1	JNICreationDelay = 10 (in seconds)
FlogiRetryCount = 3	Same as original value
FcFlogiTimeout = 10	Same as original value (in seconds)
PlogiRetryCount = 3	Same as original value
PlogiControlSeconds = 30	Same as original value (in seconds)
LunDiscoveryMethod = 1	Same as original value (LUN reporting)
CmdTaskAttr = 0	CmdTaskAttr = 0 (Simple Queue) CmdTaskAttr = 1 (Untagged)
automap = 0	automap = 1 (Enable)
FclpEnable = 1	FclpEnable = 0 (Disable)
OverrunFailoverCount = 0	Same as original value
PlogiRetryTime = 50	Same as original value
SwitchGidPtSyncEnable = 0	Same as original value
target_throttle = 256	Same as original value
lun_throttle = 64	Same as original value
Add these settings.	target0_hba = "jnic146x0"; target0_wwpn = "<controller wwpn>" target1_hba = "jnic146x1"; target1_wwpn = "<controller wwpn>"

Note: You might have to run the /etc/raid/bin/genjnicnf reconfigure script from the Solaris shell:

```
# /etc/raid/bin/genjnicnf
```

Configuration settings for FCE-1063/FCE2-1063/FCE-6410/FCE2-6410

JNI host bus adapter models FCE-1063, FCE2-1063, FCE-6410, and FCE2-6410 are supported with all currently supported levels of storage subsystem controller firmware.

Note: For each setting that is listed in Table 47, you must uncomment the line. This is true both for default settings and for settings that you must change.

Table 47. Configuration settings for FCE-1063/FCE2-1063/FCE-6410/FCE2-6410

Original value	New value
FcLoopEnabled = 1	FcLoopEnabled = 0 (for non-Loop) FcLoopEnabled = 1 (for Loop)
FcFabricEnabled = 0	FcFabricEnabled = 0 (for non-fabric) FcFabricEnabled = 1 (for fabric)
FcPortCfgEnable = 1	FcPortCfgEnable = 0 (port reconfiguration not required) FcPortCfgEnable = 1 (port reconfiguration required)
FcEngHeartbeatInterval = 5	Same as original value (in seconds)
FcLrrTimeout = 100	Same as original value (in milliseconds)
FcLinkUpRecoverTime = 1000	Same as original value (in milliseconds)
BusyRetryDelay = 5000	Same as original value (in milliseconds)
FailoverDelay = 30;	FailoverDelay = 60;
TimeoutResetEnable = 0	Same as original value
QfullRetryCount = 5	Same as original value
QfullRetryDelay = 5000	Same as original value (in milliseconds)
loRecoveryDelay = 50	Same as original value (in milliseconds)
JniCreationDelay = 5;	JniCreationDelay = 10;
FlogiRetryCount = 3	Same as original value
PlogiRetryCount = 5	Same as original value
FcEmIdEndTcbTimeCount = 1533	Same as original value
target_throttle = 256	Same as original value (default throttle for all targets)
lun_throttle = 64	Same as original value (default throttle for all LUNs)
automap = 0	automap = 0 (persistence binding) automap = 1 (automapping)
Add these settings.	target0_hba = "jnic146x0";
	target0_wwpn = "controller_wwpn"
	target1_hba = "jnic146x1";
	target1_wwpn = "controller_wwpn"

- You might have to run the /etc/raid/bin/genjnicnf reconfigure script from the Solaris shell:
/etc/raid/bin/genjnicnf
- Set portEnabled = 1; only when you see JNI cards entering non-participating mode in the /var/adm/messages file. Under that condition, complete the following steps:
 1. Set FcPortCfgEnabled = 1;
 2. Restart the host.
 3. Set FcPortCfgEnabled = 0;

4. Restart the host again.

When you have done so, check /var/adm/messages to make sure that it sets the JNI cards to Fabric or Loop mode.

Configuration settings for FCI-1063

JNI host bus adapter model FCI-1063 is supported *only* in storage subsystem configurations with controller firmware version 05.4x.xx.xx or earlier.

Note: For each setting that is listed in Table 48, you must uncomment the line. This is true both for default settings and for settings that you must change.

Table 48. Configuration settings for FCI-1063

Original value	New value
scsi_initiator_id = 0x7d	Same as original value
fca_nport = 0;	fca_nport = 1 (for the fabric) / fca_nport = 0 (for the loop)
public_loop = 0	Same as original value
target_controllers = 126	Same as original value
ip_disable = 1;	Same as original value
ip_compliant = 0	Same as original value
qfull_retry_interval = 0	Same as original value
qfull_retry_interval = 1000	Same as original value (in milliseconds)
failover = 30;	failover = 60 (in seconds)
failover_extension = 0	Same as original value
recovery_attempts - 5	Same as original value
class2_enable = 0	Same as original value
fca_heartbeat = 0	Same as original value
reset_glm = 0	Same as original value
timeout_reset_enable = 0	Same as original value
busy_retry_delay= 100;	Same as original value (in milliseconds)
link_recovery_delay = 1000;	Same as original value. (in milliseconds)
scsi_probe_delay = 500;	scsi_probe_delay = 5000 (in milliseconds; 10 milliseconds resolution)
def_hba_binding = "fca-pci*";	def_hba_binding = "nonjni"; (for binding) def_hba_binding = "fcaw"; (for non-binding)
def_wwnn_binding = "\$xxxxxx"	def_wwnn_binding = "xxxxxx"
def_wwpn_binding = "\$xxxxxx"	Same as the original entry
fca_verbose = 1	Same as the original entry
Will be added by reconfigure script	name="fca-pci" parent="physical path" unit-address="#"
Will be added by reconfigure script	target0_hba="fca-pci0" target0_wwpn="controller wwpn";
Will be added by reconfigure script	name="fca-pci" parent="physical path" unit-address="#"
Will be added by reconfigure script	target0_hba="fca-pci1" target0_wwpn= "controller wwpn";

Note: You might have to run the /etc/raid/bin/genjnicnf reconfigure script from the Solaris shell:

```
# /etc/raid/bin/genjnicnf
```

Configuration settings for FC64-1063

JNI host bus adapter model FC64-1063 is supported *only* in storage subsystem configurations with controller firmware version 05.4x.xx.xx, or earlier.

Important: For each setting that is listed in Table 49, you must uncomment the line. This is true both for default settings and for settings that you must change.

Table 49. Configuration settings for FC64-1063

Original value	New value
fca_nport = 0;	fca_nport =1;
ip_disable = 0;	ip_disable=1;
failover = 0;	failover =30;
busy_retry_delay = 5000;	busy_retry_delay = 5000;
link_recovery_delay = 1000;	link_recovery_delay = 1000;
scsi_probe_delay = 5000;	scsi_probe_delay = 5000;
def_hba_binding = "fcaw*";	Direct attached configurations: def_hba_binding = "fcaw*"; SAN-attached configurations: def_hba_binding = "nonJNI";
def_wwnn_binding = "\$xxxxxx"	def_wwnn_bindindef_hba_ binding = "nonjni"; g = "xxxxxx"
def_wwnn_binding = "\$xxxxxx"	Same as the original entry
Will be added by reconfigure script	name="fcaw" parent="<physical path>" unit-address="<#>"
Will be added by reconfigure script	target0_hba="fcaw0" target0_wwpn="<controller wwpn>";
Will be added by reconfigure script	name="fcaw" parent="<physical path>" unit-address="<#>"
Will be added by reconfigure script	target0_hba="fcaw0" target0_wwpn= "<controller wwpn>";

Note: You might have to run the /etc/raid/bin/genscsiconf reconfigure script from the shell prompt:

```
# /etc/raid/bin/genscsiconf
```

QLogic HBA settings

The QLogic HBAs are not Plug and Play with autoconfiguration. Instead, you must change the settings or bindings, as described in Table 50 on page 284.

Note: In Table 50 on page 284, the HBA is identified as hba0. However, you must modify the settings on both QLogic HBAs: hba0 and hba1.

When you modify the settings on hba1, use the same values that are listed in the table, but change all instances of hba0 to hba1, as shown in the following example.

HBA	Original value	New value
hba0	hba0-execution-throttle=16;	hba0-execution-throttle=255;
hba1	hba1-execution-throttle=16;	hba1-execution-throttle=255;

In the vi Editor, uncomment and modify the loop attributes of each QLogic HBA, using the values that are specified in Table 50.

Table 50. Configuration settings for QL2342

Original value	New value	Comments
max-frame-length=2048;	max-frame-length=2048	Use the default
execution-throttle=16;	execution-throttle=255;	Change
login-retry-count=8;	login-retry-count=30;	Change
enable-adapter-hard-loop-ID=0;	enable-adapter-hard-loop-ID=1;	Change
adapter-hard-loop-ID=0;	adapter-hard-loop-ID=0;	Must be a unique number
enable-LIP-reset=0;	enable-LIP-reset=0;	Use the default
hba0-enable-LIP-full-login=1;	hba0-enable-LIP-full-login=1;	Use the default
enable-target-reset=0;	enable-target-reset=0;	Use the default
reset-delay=5	reset-delay=8	Change
port-down-retry-count=8;	port-down-retry-count=70;	Change
maximum-luns-per-target=8;	maximum-luns-per-target=0;	Change
connection-options=2;	connection-options=2;	Use the default
fc-tape=1;	fc-tape=0;	Change
loop-reset-delay = 5;	loop-reset-delay = 8;	Change
> gbyte-addressing = disabled;	> gbyte-addressing = enabled;	Change
link-down-timeout = 30;	link-down-timeout = 60;	Change

Appendix B. Using a storage subsystem with a VMware ESX Server configuration

The Storage Manager software is not currently available for VMware ESX Server operating systems. Therefore, to manage DS3000, DS4000, and DS5000 storage subsystems with your VMware ESX Server host, you must install the Storage Manager client software (SMclient) on a Windows or Linux management station. This can be the same workstation that you use for the browser-based VMware ESX Server Management Interface. Also, to enable Asymmetric Logical Unit Access (ALUA) you must have VMware ESX Server operating system version 4.1 u2 or later and 5.0 u1 or later.

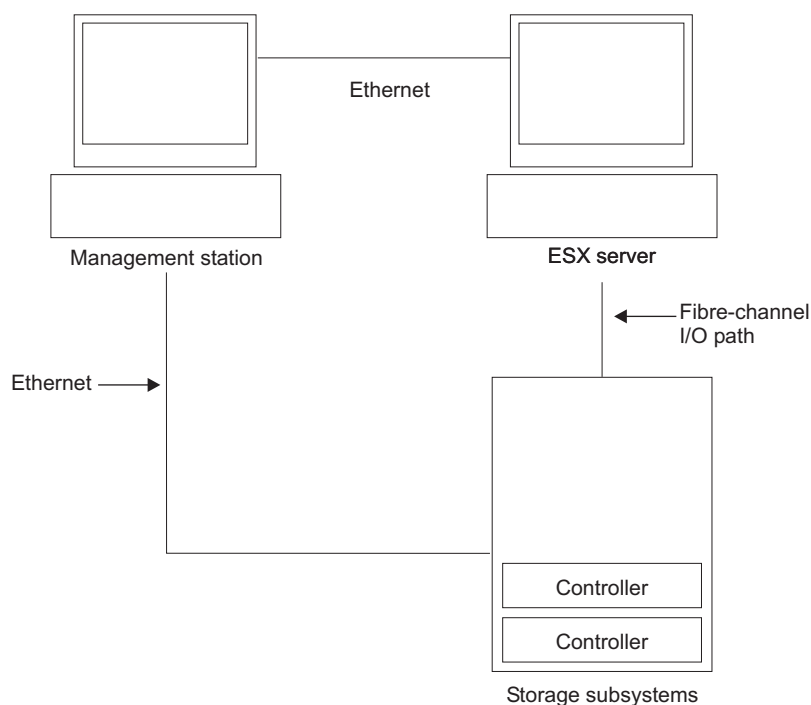
For additional information about using a DS3000, DS4000, or DS5000 storage subsystem with a VMware ESX Server host, see “VMware ESX Server restrictions” on page 287.

You can also see the System Storage Interoperation Center at the following website:

www.ibm.com/systems/support/storage/config/ssic

Sample configuration

Figure 34 shows a sample VMware ESX Server configuration.



SJ001150

Figure 34. Sample VMware ESX Server configuration

Software requirements

This section describes the software that is required to use a VMware ESX Server host operating system with a DS3000, DS4000, or DS5000 storage subsystem.

Management station

The following software is required for the Windows or Linux management station:

1. SM Runtime (Linux only)
2. SMclient (Linux and Windows)

Host (VMware ESX Server)

The following software is required for VMware ESX Server:

- VMware ESX Server (with DS3000, DS4000, or DS5000 controller firmware version 07.1x.xx.xx)
- VMware ESX Server-supplied driver for the Fibre Channel HBAs
- VMware ESX Server-supplied QLogic driver failover setup
- VMware ESX Server Tools (installed on all virtual machines using DS3000, DS4000 or DS5000 logical drives)

Earlier versions of VMware ESX Server:

1. VMware ESX Server 2.1 was supported with DS4000 and DS5000 controller firmware version 06.12.xx.xx only.
2. VMware ESX Server 2.0 was supported with DS4000 and DS5000 controller firmware version 05.xx.xx.xx only.

Guest OS Clustering: If you intend to create a Guest OS cluster configuration, you must use Microsoft Cluster Services software, in addition to the host software requirements listed in this section.

VMware Host Clustering: VMware ESX Server 2.5 and higher comes with a Distributed Resource Scheduler and high availability for clustering, which allows you to aggregate several hosts' resources into one resource pool. (A DRS cluster is implicitly a resource pool.)

For information about Windows clustering with VMware ESX Server, see the ESX Server 2.5 Installation Guide at the following website: <http://www.vmware.com/support/pubs/>.

Hardware requirements

You can use VMware ESX Server host servers with the following types of storage subsystems and storage expansion enclosures. For additional information, see the System Storage Interoperation Center at the following website:

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

Note: For general storage subsystem requirements, see Chapter 1, "Preparing for installation," on page 1.

DS5000 Storage Subsystems

- DS5300
- DS5100

DS4000 Storage Subsystems

- DS4100 (Dual-controller units only)
- DS4200
- DS4300 (Dual-controller and Turbo units only)
- DS4400
- DS4500
- DS4700
- DS4800

DS5000 storage expansion enclosures

- EXP5000

DS4000 storage expansion enclosures

- EXP100
- EXP420 (with DS4200 only)
- EXP500
- EXP700
- EXP710
- EXP810

VMware ESX Server restrictions

SAN and connectivity restrictions:

- VMware ESX Server hosts support host-agent (out-of-band) managed storage subsystem configurations only. Direct-attached (in-band) management configurations are not supported.
- VMware ESX Server hosts can support multiple host bus adapters (HBAs) and DS3000, DS4000, and DS5000 devices. However, there is a restriction on the number of HBAs that can be connected to a single storage subsystem. You can configure up to two HBAs per partition and up to two partitions per storage subsystem. Additional HBAs can be added for additional storage subsystems and other SAN devices, up to the limits of your specific storage subsystem platform.
- When you are using two HBAs in one VMware ESX Server, LUN numbers must be the same for each HBA attached to the storage subsystem.
- Single HBA configurations are allowed, but each single HBA configuration requires that both controllers in the storage subsystem be connected to the HBA through a switch. If they are connected through a switch, both controllers must be within the same SAN zone as the HBA.

Attention: A single HBA configuration might result in loss of data access in the event of a path failure.

- Single-switch configurations are allowed, but each HBA and storage subsystem controller combination must be in a separate SAN zone.
- Other storage devices, such as tape devices or other disk storage, must be connected through separate HBAs and SAN zones.

Partitioning restrictions:

- The maximum number of partitions per VMware ESX Server host, per storage subsystem, is two.

- All logical drives that are configured for VMware ESX Server must be mapped to an VMware ESX Server host group.

Note: VMware ESX server-specific host type is not available for DS3000, DS4000, or DS5000 storage subsystems if the controller firmware version is earlier than 7.70.xx.xx. Use LNXCLVMWARE host type for your VMware hosts and host groups. If you are using the default host group, make sure that the default host type is LNXCLVMWARE. DS Storage subsystems with controller firmware version 7.70.xx.xx or later have a VMware ESX server-specific host type defined, named VMWARE. VMWARE should be used as the host type of the VMWare hosts and host group.

- In a DS4100 storage subsystem configuration, you must initially assign the LUNs to Controller A, on the lowest-numbered HBA. After the LUNs are initialized, you can change the path to Controller B. (This restriction will be corrected in a future release of ESX Server.)
- Assign LUNs to the ESX Server starting with LUN number 0.
- Do not map an access (UTM) LUN to any of the ESX Server hosts or host groups. Access (UTM) LUNs are used only with in-band managed storage subsystem configurations, which VMware ESX Server does not support at this time.

Failover restrictions:

- You must use the VMware ESX Server failover driver for multipath configurations. Other failover drivers (such as RDAC) are not supported in VMware ESX Server configurations.
- The default failover policy for all storage subsystems is now MRU (most recently used).
- Use the LNXCLVMWARE (if the controller firmware is earlier than 7.70.xx.xx) or VMWARE (if the controller firmware is 7.70.xx.xx or later) host type in VMware ESX Server configurations (2.0 and higher). The LNXCLVMWARE or VMWARE host type automatically disables Auto Drive Transfer (ADT).

Interoperability restrictions:

- DS4100 and DS4300 *single-controller* storage subsystems are not supported with VMware ESX Server hosts. (DS4100 and DS4300 *dual-controller* storage subsystems are supported.)
- EXP700 storage expansion enclosures are not supported with DS4800 storage subsystems. You must upgrade to EXP710 storage expansion enclosures.

Other restrictions:

- Dynamic Logical Drive Expansion (DVE) is not supported for VMFS-formatted LUNs on VMware ESX server operating system earlier than 2.5.x. For information about availability of DS Copy Service features that are supported VMware ESX Server 2.5 Server and higher configurations, contact your IBM support representative.
- Do not boot your system from a SATA device.

Other VMware ESX Server host information

For more information about setting up your VMware ESX Server host, see the documentation and readme files that are maintained at the following website:

www.vmware.com/support/pubs/

For information about installing a VMware ESX Server operating system on an IBM server, see the IBM support website at:

www-03.ibm.com/systems/i/advantages/integratedserver/vmware/

Configuring storage subsystems for VMware ESX Server

Before you can configure a storage subsystem, you must physically configure the host server, SAN fabric, and storage subsystem controllers; assign initial IP addresses to the controllers; and install SMclient on the Windows or Linux management station. See Chapter 4, “Configuring storage,” on page 55 for storage subsystem configuration procedures.

Cross-connect configuration for VMware connections

A cross-connect Storage Area Network (SAN) configuration is required when VMware hosts are connected to a DS3000, DS4000, or DS5000 storage subsystem. Each Host Bus Adapter (HBA) in a VMware host must have a path to each of the controllers in the storage subsystem. Figure 35 shows the cross connections for VMware server configurations.




Figure 35. Cross-connect configuration for VMware connections

Mapping LUNs to a storage partition on VMware ESX Server

See “Mapping LUNs” on page 84 for procedures that describe how to map the LUNs to a partition. This section contains notes about LUN mapping that are specific to VMware ESX Servers.

When you map your LUNs on VMware ESX Server, note the following:

- Map the LUNs using consecutive numbers, starting with LUN 0. For example, map LUNs to numbers 0; 1; 2; 3; 4; 5; and so on, without skipping any numbers.
- On each partition, you must map a LUN 0.
- If your configuration does not require LUN sharing (single or multiple independent ESX Servers, local virtual cluster), each logical drive must be mapped either directly to a host, or to a host group with a single host as a member.
- LUN sharing across multiple ESX servers is only supported when you are configuring VMotion enabled hosts or Microsoft Cluster nodes. On LUNs that are mapped to multiple ESX Servers, you must change the access mode to Shared.

You can map the LUNs to a host group for the ESX Servers, so they will be available to all members of the host group. For additional information on Windows Clustering with ESX Server, see the *ESX Installation Guide* at the following website:

www.vmware.com/support/pubs/

Verifying the storage configuration for VMware

Complete the following steps to verify that your storage subsystem is set up correctly and that you can see the storage subsystem:

1. Start the server.
2. After QLogic BIOS initialization, press Ctrl+Q to enter the Fast!UTIL setup program.
3. Select the first host bus adapter that is displayed on the Fast!UTIL screen.
4. Select **Host Adapter Settings**, and press Enter.
5. Select **Scan Fibre Devices** and press Enter. The resulting output is similar to the following:

ID	Vendor	Product	Rev	Port Name	Port ID
128	No device present		0520		
129	IBM	1742	0520	200400A0b00F0A16	610C00
130	No device present				
131	No device present				
132	No device present				
133	No device present				
134	No device present				
135	No device present				

Note: Depending on how the configuration is cabled, you might see multiple instances.

If you do not see a storage subsystem controller, verify the cabling, switch zoning, and LUN mapping.

Appendix C. Using the Storage Manager with high-availability cluster services

The high-availability clustering services provided by the Storage Manager allow application services to continue when a hardware or software failure occurs. This system protects you from software failures as well as from the failure of a CPU, disk, or LAN component. If a component fails, its redundant partner component takes over cluster services and coordinates the transfer between components.

General information

This document does not describe how to install or configure cluster services. See the documentation that is provided with your cluster service products for this information.

Important: The information in this document might not include up-to-date cluster software version levels.

For the latest requirements and user information about using the Storage Manager with cluster services, see the readme file that is located on the Storage Manager DVD for your host operating system, or check the most recent readme files online.

See “Finding Storage Manager software, controller firmware, and readme files” on page xiv for instructions on finding the readme files online.

You can also find more information on the System Storage Interoperation Center, which is maintained at the following website:

www.ibm.com/systems/support/storage/config/ssic

Using cluster services on AIX systems

The following sections contain general hardware requirements and additional information about cluster services.

Important: The information in this document might not show up-to-date cluster software version levels. Check the Storage Manager readme file for AIX for up-to-date information about clustering requirements. See “Finding Storage Manager software, controller firmware, and readme files” on page xiv for instructions on finding the readme file on the web.

You can also see the following web sites for the most current information about AIX and clustering:

www.ibm.com/systems/support/storage/config/ssic

publib.boulder.ibm.com/infocenter/clresctr/index.jsp

High-Availability Cluster Multi-Processing

This section contains general requirements and usage notes for High Availability Cluster Multi-Processing (HACMP™) support with the Storage Manager.

Software requirements

For the latest supported HACMP versions, see the System Storage Interoperation Center at the following website:

www.ibm.com/systems/support/storage/config/ssic

Configuration limitations

The following limitations apply to HACMP configurations:

- HACMP C-SPOC cannot be used to add a DS3000, DS4000, or DS5000 disk to AIX using the *Add a Disk to the Cluster* facility.
- HACMP C-SPOC does not support enhanced concurrent mode arrays.
- Single-HBA configurations are allowed, but each single-HBA configuration requires that both controllers in the storage subsystem be connected to a switch, within the same SAN zone as the HBA.

Attention: Although single-HBA configurations are supported, do not use them in HACMP environments because they introduce a single point-of-failure in the storage I/O path.

- Use switched fabric connections between the host nodes and the storage subsystem. Direct attachment from the host nodes to the storage subsystem in an HACMP environment is supported *only* if all the following restrictions and limitations are met:
 - Only dual-controller DS3000, DS4000, or DS5000 storage subsystem versions are supported for direct attachment in a high-availability configuration.
 - The AIX operating system must be version 05.2 or later.
 - The HACMP clustering software must be version 05.1 or later.
 - All host nodes that are directly attached to the storage subsystem must be part of the same HACMP cluster.
 - All logical drives (LUNs) that are surfaced by the storage subsystem are part of one or more enhanced concurrent mode arrays.
 - The array **varyon** is in the active state *only* on the host node that owns the HACMP non-concurrent resource group (which contains the enhanced concurrent mode array or arrays). For all other host nodes in the HACMP cluster, the enhanced concurrent mode array **varyon** is in the passive state.
 - Direct operations on the logical drives in the enhanced concurrent mode arrays cannot be performed, from any host nodes in the HACMP cluster, if the operations bypass the Logical VolumeManager (LVM) layer of the AIX operating system. For example, you cannot use a DD command while logged in as the root user.
 - Each host node in the HACMP cluster must have two Fibre Channel connections to the storage subsystem. One direct Fibre Channel connection must be to controller A in the storage subsystem, and the other direct Fibre Channel connection must be to controller B in the storage subsystem.
 - You can directly attach a maximum of two host nodes in an HACMP cluster to a dual-controller version of a DS4100 or DS4300 storage subsystem.
 - You can directly attach a maximum of two host nodes in an HACMP cluster to a storage subsystem. Each host node must have two direct Fibre Channel connections to the storage subsystem.

Note: In a DS3000, DS4000, or DS5000 storage subsystem, the two direct Fibre Channel connections from each host node must be to independent minihubs. Therefore, this configuration requires that four host minihubs (feature code

3507) be installed in the DS3000, DS4000, or DS5000 storage subsystem—two host minihubs for each host node in the HACMP cluster.

Other HACMP usage notes

The following notations are specific to HACMP environments:

- HACMP clusters can support from two to 32 servers on each DS3000, DS4000, and DS5000 storage subsystem partition. If you run this kind of environment, be sure to read and understand the AIX device drivers queue depth settings that are described in “Setting the queue depth for hdisk devices” on page 164.
- You can attach non-clustered AIX hosts to a storage subsystem that is running the Storage Manager and is attached to an HACMP cluster. However, you must configure the non-clustered AIX hosts on separate host partitions on the storage subsystem.

Parallel System Support Programs and General Parallel File System

This section contains general requirements and usage notes for Parallel System Support Programs (PSSP) and General Parallel File System (GPFS™) support with the Storage Manager.

Software requirements

For the latest supported PSSP and GPFS versions, see the System Storage Interoperation Center at the following website:

www.ibm.com/systems/support/storage/config/ssic

Configuration limitations

The following limitations apply to PSSP and GPFS configurations:

- Direct connection is not allowed between the host node and a DS3000, DS4000, or DS5000 storage subsystem. Only switched fabric connection is allowed.
- RVSD clusters can support up to two IBM Virtual Shared Disk and RVSD servers for each storage subsystem partition.
- Single node quorum is not supported in a dual-node GPFS cluster with DS3000, DS4000, or DS5000 disks in the configuration.
- Heterogeneous configurations are not supported.

Other PSSP and GPFS usage notes

In GPFS file systems, the following DS3000, DS4000, and DS5000 storage subsystem cache settings are supported:

- Read cache enabled or disabled
- Write cache enabled or disabled
- Cache mirroring enabled or disabled (depending upon the write cache mirroring setting)

The performance benefits of read or write caching depends on the application.

GPFS, PSSP, and HACMP cluster configuration diagrams

The diagrams in this section show both the preferred and failover paths from an HBA pair to a given logical drive or set of logical drives.

A preferred path to a logical drive is determined when the logical drive is created and distributed across a storage subsystem controller. The controller to which it is assigned determines which path is preferred or active for I/O transfer. Logical

drives can, and in most cases must, be assigned to both controllers, balancing the I/O load across HBAs and storage subsystem controllers.

Figure 36 shows a cluster configuration that contains a single DS storage subsystem, with one to four partitions.




Figure 36. Cluster configuration with single storage subsystem—one to four partitions

Figure 37 on page 295 shows a cluster configuration that contains three DS storage subsystems, with one partition on each storage subsystem.

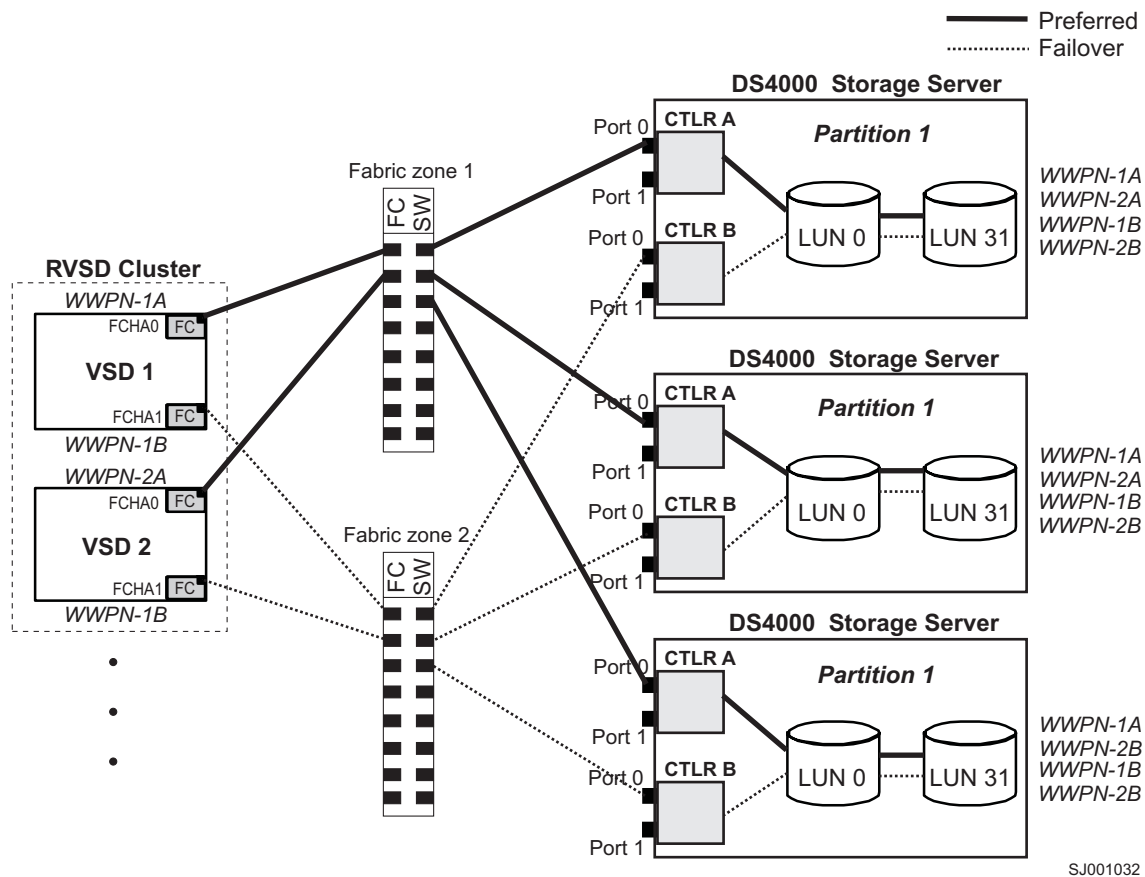



Figure 37. Cluster configuration with three storage subsystems—one partition per subsystem

Figure 38 on page 296 shows a cluster configuration that contains four DS storage subsystems, with one partition on each storage subsystem.



SJ001033

Figure 38. Cluster configuration with four storage subsystems—one partition per subsystem

Figure 39 on page 297 shows a cluster configuration that contains two DS storage subsystems, with two partitions on each storage subsystem.




Figure 39. RVSD cluster configuration with two storage subsystems—two partitions per subsystem

Figure 40 on page 298 shows an HACMP/GPFS cluster configuration that contains a single DS storage subsystem, with one partition.




Figure 40. HACMP/GPFS cluster configuration with one storage subsystem—one partition

Figure 41 on page 299 shows an HACMP/GPFS cluster configuration that contains two DS storage subsystems, with two partitions on each storage subsystem.

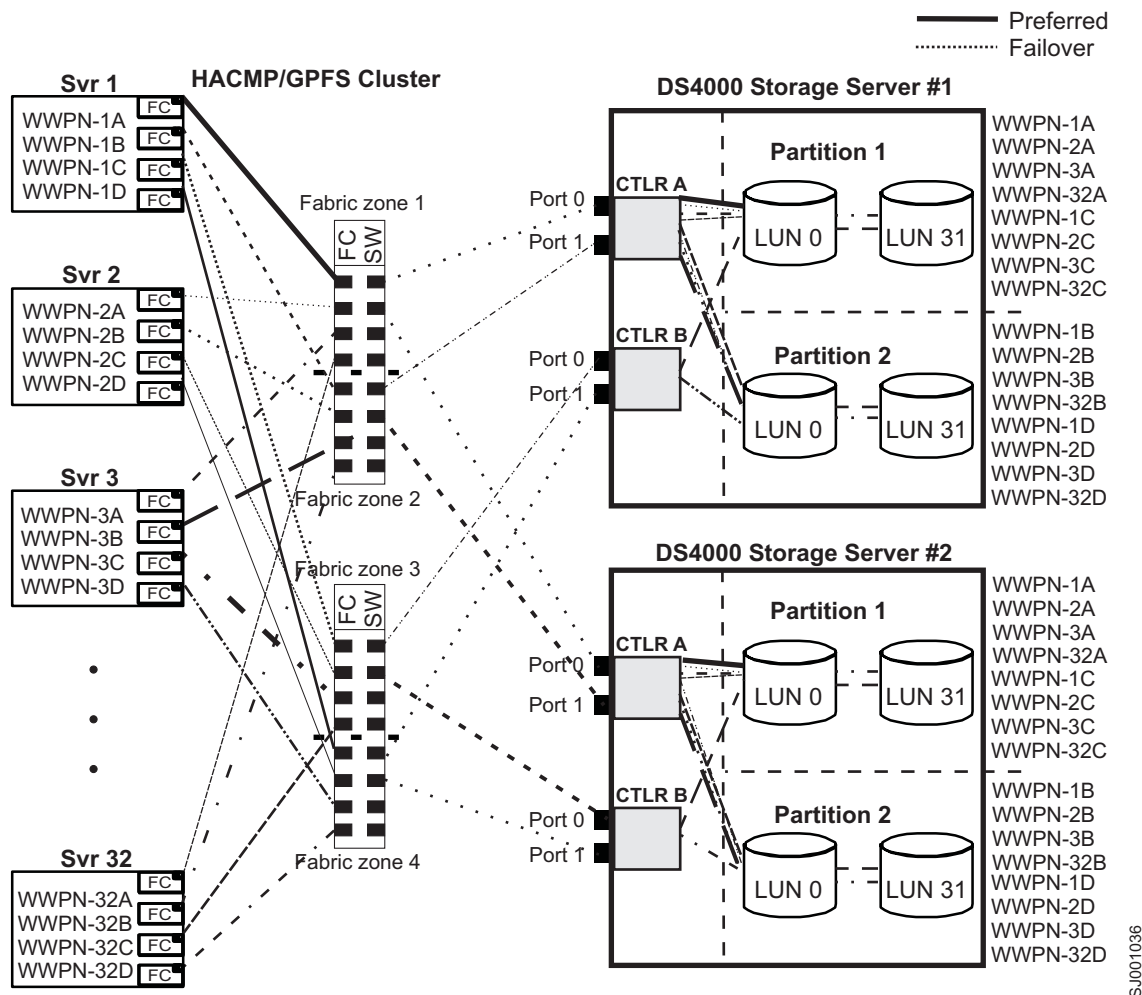


Figure 41. HACMP/GPFS cluster configuration with two storage subsystems—two partitions per subsystem

Using cluster services on HP-UX systems

The information in this document might not show up-to-date cluster software version levels. Check the Storage Manager readme file for HP-UX for up-to-date information about clustering requirements. See “Finding Storage Manager software, controller firmware, and readme files” on page xiv for instructions on finding the readme file online.

You can also see the System Storage Interoperation Center at the following website:

www.ibm.com/systems/support/storage/config/ssic

You can choose among many configurations when you set up clustering on an HP-UX system. A minimum configuration consists of two servers that are configured with both a primary and two standby LANs to establish a heartbeat LAN.

Provide Fibre Channel connections to the storage subsystem through two switches that provide the necessary redundant data path for the hosts. Make sure that each server has two HP Tachyon host bus adapters.

Using cluster services on Solaris systems

The following sections contain general hardware requirements and additional information about cluster services.

Important: The information in this document might not show up-to-date cluster software version levels. Check the Storage Manager readme file for Solaris for up-to-date information about clustering requirements, including the latest supported versions of Veritas Cluster Server. See “Finding Storage Manager software, controller firmware, and readme files” on page xiv for instructions on finding the readme file online.

You can also see the System Storage Interoperation Center at the following website:

www.ibm.com/systems/support/storage/config/ssic

General Solaris requirements

Each Solaris system in the cluster requires the following hardware:

- At least three Ethernet ports:
 - Two for the private network connections
 - At least one for the public network connection
- Two Fibre Channel host bus adapters for connection to the storage subsystem
- A SCSI connection for operating system disks
- Each Veritas Cluster Server system requires at least 128 MB of RAM and 35 MB of free disk space

System dependencies

This section provides information about RDAC IDs and single points of failure.

Adding RDAC IDs

You can add up to eight additional IDs to the `/etc/symsm/rmparams` file. Complete the following steps to add them:

1. Open the `/etc/symsm/rmparams` file in the vi Editor by typing the following command:

```
# vi /etc/symsm/rmparams
```
2. Modify the `Rdac_HotAddIDs` line as follows:

```
Rdac_HotAddIDs:0:1:2:3:4:5:6:7:8
```
3. Save and close the `/etc/symsm/rmparams` file.

Single points of failure

When setting up cluster services, it is important to eliminate single points of failure because a single point of failure makes a cluster only as strong as its weakest component. Set up the storage subsystem for shared storage; for example, all the nodes in the cluster must recognize the same storage and the host types must be set correctly.

Appendix D. Viewing and setting AIX Object Data Manager (ODM) attributes

Some of the ODM attributes are for information purposes only. These information-only attributes show how the storage subsystem is configured or its current state. You can modify other attributes using SMIT or with the UNIX **chdev -p** command.

Attribute definitions

The following tables list definitions and values of the ODM attributes for dars, dacs and hdisks:

- Table 51: *Attributes for dar devices*
- Table 52 on page 303: *Attributes for dac devices*
- Table 53 on page 303: *Attributes for hdisk devices*

Note:

1. Attributes with True in the Changeable column can be modified from their default settings.
2. Attributes with False in the Changeable column are for informational or state purposes only. However, some attributes with False in the Changeable column can be modified using the Storage Manager.
3. The **lsattr -El** (uppercase E, lowercase L) command is another way to determine which attributes can be modified. Attributes that can be modified display True in the last column of the **lsattr -El** output. You can also display the default values with the **lsattr -Dl** command.

Table 51. *Attributes for dar devices*

Attribute	Definition	Changeable (T/F)	Possible value
<i>act_controller</i>	List of controllers in the active state at the time of configuration.	False	Set at configuration time by the RDAC software.
<i>all_controller</i>	List of controllers that comprise this array; usually there are two dac devices.	False	Set at configuration time by the RDAC software.
<i>held_in_reset</i>	Name of the controller that was in the held-in-reset state at the time of configuration, or none if no controllers were in that state.	True	Set at configuration time by the RDAC software. Must not be changed.
<i>load_balancing</i>	Indicator that shows whether load balancing is enabled (yes) or disabled (no); see the definition of the <i>balance_freq</i> attribute for more information.	True	Yes or No. Attention: You must only set the <i>load_balancing</i> attribute to yes in single-host configurations.

Table 51. Attributes for dar devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<i>autorecovery</i>	Indicator that shows whether the device returns the array to dual-active mode when it detects proper operation of both paths and controllers (yes) or not (no).	True	Yes or No. See restrictions on use.
<i>hlthchk_freq</i>	Number that specifies how often health checks are performed, in seconds.	True	1 - 9999. Must not be changed.
<i>aen_freq</i>	Number that specifies how often polled AEN checks are performed, in seconds.	True	1 - 9999. Must not be changed.
<i>balance_freq</i>	If <i>load_balancing</i> is enabled, number that specifies how often the system performs load-balancing on the array, in seconds.	True	1 - 9999 - Must not be changed.
<i>fast_write_ok</i>	Indicator that shows whether fast-write write-caching is available for this system (yes) or not (no).	False	Yes or No. State of the storage subsystem configuration.
<i>cache_size</i>	Cache size for both controllers, in megabytes; 0 if the sizes do not match.	False	512 or 1024. Set by the storage subsystem.
<i>switch_retries</i>	Number that specifies how many times to retry failed switches, in integers.	True	0 - 255. Default: 5 For most configurations, the default is the best setting. If you are using HACMP, it can be helpful to set the value to 0. Attention: You cannot use concurrent firmware download if you change the default setting.

Table 52. Attributes for dac devices

Attribute	Definition	Changeable (T/F)	Possible value
<i>passive_control</i>	Indicator that shows whether this controller was in passive state at the time of configuration (yes) or not (no).	False	Yes or No. State of the storage subsystem configuration.
<i>alt_held_reset</i>	Indicator that shows whether the alternate controller was in the held-in-reset state at the time of configuration (yes) or not (no).	False	Yes or No. State of the storage subsystem configuration.
<i>controller_SN</i>	Serial number of this controller.	False	Set by the storage subsystem.
<i>ctrl_type</i>	Type of array to which this controller belongs.	False	1742, 1722, 1742-900. Set by the storage subsystem.
<i>cache_size</i>	Cache size of this controller, in megabytes.	False	512, 1024. Set by the storage subsystem.
<i>scsi_id</i>	SCSI identifier of this controller.	False	Set by SAN, reported by AIX.
<i>lun_id</i>	Logical unit number of this controller.	False	Set by the storage subsystem.
<i>utm_lun_id</i>	Logical unit number of this controller, or none if UTM (access logical drives) is not enabled.	False	0 - 31. Set by the Storage Manager.
<i>node_name</i>	Name of the Fibre Channel node.	False	Set by the storage subsystem.
<i>location</i>	User-defined location label for this controller; the system does not use this value.	True	Set by the Storage Manager.
<i>ww_name</i>	Fibre Channel worldwide name of this controller.	False	Set by the storage subsystem.
<i>GLM_type</i>	GLM type used for this controller.	False	High or Low. Set by the storage subsystem.

Table 53. Attributes for hdisk devices

Attribute	Definition	Changeable (T/F)	Possible value
<i>pvid</i>	AIX physical volume identifier, or none if not set.	False	Set by AIX.

Table 53. Attributes for hdisk devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<i>q_type</i>	Queueing type for this device; must be set to simple .	False	Set by AIX. Must be "simple".
<i>queue_depth</i>	Number that specifies the depth of the queue based on system configuration; reduce this number if the array is returning a BUSY status on a consistent basis.	True	1 - 64 Note: See "Setting the queue depth for hdisk devices" on page 164 for important information about setting this attribute.
<i>PR_key_value</i>	Required only if the device supports any of the persistent reserve policies. This attribute is used to distinguish between different hosts.	True	1-64, or None. Note: You must set this attribute to non-zero before the reserve_policy attribute is set.
<i>reserve_policy</i>	Persistent reserve policy, which defines whether a reservation methodology is employed when the device is opened.	True	no_reserve PR_shared, PR_exclusive, or single_path
<i>max_transfer</i>	Maximum transfer size is the largest transfer size that can be used in sending I/O.	True	Numeric value; Default = 1 MB Note: Usually unnecessary to change default, unless very large I/Os require increasing the value.
<i>write_cache</i>	Indicator that shows whether write-caching is enabled on this device (yes) or not (no); see the definition of the <i>cache_method</i> attribute for more information.	False	Yes or No.
<i>size</i>	Size of this logical drive.	False	Set by the storage subsystem.
<i>raid_level</i>	Number that specifies the RAID level of this device.	False	0, 1, 3, 5. Set by the Storage Manager.
<i>rw_timeout</i>	Number that specifies the read/write timeout value for each read/write command to this array, in seconds; usually set to 30.	True	30 - 180. Must not be changed from default.

Table 53. Attributes for hdisk devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<i>reassign_to</i>	Number that specifies the timeout value for FC reassign operations, in seconds; usually set to 120.	True	0 - 1000. Must not be changed from default.
<i>scsi_id</i>	SCSI identifier at the time of configuration.	False	Set by SAN, reported by AIX.
<i>lun_id</i>	Logical unit number of this device.	False	0 - 255. Set by the Storage Manager.
<i>cache_method</i>	<p>If <i>write_cache</i> is enabled, the write-caching method of this array; set to one of the following:</p> <ul style="list-style-type: none"> • default. Default mode; the word "default" is not seen if <i>write_cache</i> is set to yes. • fast_write. Fast-write (battery-backed, mirrored write-cache) mode. • fw_unavail. Fast-write mode was specified but could not be enabled; write-caching is not in use. • fast_load. Fast-load (non-battery-backed, non-mirrored write-cache) mode. • fl_unavail. Fast-load mode was specified but could not be enabled. 	False	Default, fast_write, fast_load, fw_unavail, fl_unavail.
<i>prefetch_mult</i>	Number of blocks to be prefetched into read cache for each block read.	False	0 - 100.
<i>ieee_volname</i>	IEEE unique logical drive name identifier for this logical drive.	False	Set by the storage subsystem.

Using the lsattr command to view ODM attributes

To view the Object Data Manager (ODM) attribute settings for dars, dacs, and hdisks, use the **lsattr** command, as follows:

- To view the default settings, type **lsattr -Dl**.
- To view the attributes that are currently set on the system, type **lsattr -El**.

The **lsattr -El** output examples shown in Table 54, Table 55, and Table 56 on page 307, display the ODM attribute settings for a dar, a dac and an hdisk.

Table 54. Example 1: Displaying the attribute settings for a dar

```
# lsattr -El dar0
act_controller dac0,dac1 Active Controllers          False
aen_freq      600      Polled AEN frequency in seconds      True
all_controller dac0,dac1 Available Controllers          False
autorecovery  no       Autorecover after failure is corrected      True
balance_freq  600      Dynamic Load Balancing frequency in seconds      True
cache_size    128      Cache size for both controllers          False
fast_write_ok yes      Fast Write available          False
held_in_reset none     Held-in-reset controller          True
hlthchk_freq  600      Health check frequency in seconds      True
load_balancing no       Dynamic Load Balancing          True
switch_retries 5        Number of times to retry failed switches      True
```

Table 55. Example 2: Displaying the attribute settings for a dac

```
# lsattr -El dac0
GLM_type      low      GLM type          False
alt_held_reset no       Alternate held in reset      False
cache_size    128      Cache Size in MBytes        False
controller_SN 1T24594458      Controller serial number      False
ctrl_type     1722-600   Controller Type          False
location      location      Location Label          True
lun_id        0x0      Logical Unit Number        False
node_name     0x200200a0b80f14af FC Node Name          False
passive_control no       Passive controller          False
scsi_id       0x11000      SCSI ID                  False
utm_lun_id    0x001f000000000000 Logical Unit Number        False
ww_name       0x200200a0b80f14b0 World Wide Name          False
```

Note: Running the **# lsattr -Rl <device> -a <attribute>** command, will show allowable values for the specified attribute and is an hdisk attribute list when using MPIO.

Note: In Table 56 on page 307, the **ieee_volname** and **lun_id** attribute values are shown abbreviated. An actual output would show the values in their entirety.

Table 56. Example 3: Displaying the attribute settings for an hdisk

lsattr -El hdisk174			
cache_method	fast_write	Write Caching method	False
ieee_volname	600A0B8...1063F7076A7	IEEE Unique volume name	False
lun_id	0x0069...000000	Logical Unit Number	False
prefetch_mult	12	Multiple of blocks to prefetch on read	False
pvid	none	Physical volume identifier	False
q_type	simple	Queuing Type	False
queue_depth	2	Queue Depth	True
raid_level	5	RAID Level	False
reassign_to	120	Reassign Timeout value	True
reserve_lock	yes	RESERVE device on open	True
rw_timeout	30	Read/Write Timeout value	True
scsi_id	0x11f00	SCSI ID	False
size	2048	Size in Mbytes	False
write_cache	yes	Write Caching enabled	False

Appendix E. About VDS/VSS provider

The Microsoft Virtual Disk Service (VDS) and Microsoft Volume Shadow-copy Service (VSS) are IBM DS Storage Manager interfaces for Microsoft Windows Server 2003 and Microsoft Windows Server 2008. VDS and VSS enable the storage subsystem to interact with third-party applications that use the VDS or VSS Application Programming Interface (API). Microsoft VDS/VSS is included in the Windows Server 2003 and Windows Server 2008 installations. Microsoft VSS is included and supported in Windows Server 2012, but Microsoft VDS is deprecated from Windows Server 2012.

Note: IBM VDS/VSS hardware provider does not support client versions of Windows Server 2012.

IBM VDS/VSS hardware provider does not support client versions of Windows Server 2012. The IBM DS Storage Manager VDS hardware provider is a Windows Dynamic Link Library (DLL) that VDS loads, and is used as a communication channel to the storage subsystem. With the IBM DS Storage Manager VDS hardware provider installed, third-party applications can send management commands to the storage subsystem. It supports commands like - creating logical drives, deleting logical drives, and unmasking logical drives. Third-party applications can obtain status and configuration information from the storage subsystem. The IBM DS Storage Manager VSS hardware provider is a Windows service (.exe). Microsoft's VSS attaches to the service and uses it to coordinate the creation of FlashCopy logical drives on the storage subsystem. VSS-initiated logical drive FlashCopies can be triggered through third-party backup tools known as 'requestors'.

Download the VDS/VSS provider and the installation instructions from the following website. First, register a free account and obtain login credentials and then download the installer.

<http://support.netapp.com/NOW/apbu/oemcp/protcd/>.

Appendix F. Installing SMI-S provider

SMIS-S (Storage Management Initiative Specification) defines a method for interoperable management of a heterogeneous Storage Area Network (SAN), and describes the information available to a WBEM Client from an SMI-S compliant CIM Server and an object-oriented, XML-based, messaging-based interface. This interface is designed to support specific requirements of managing devices in and through SANs, IBM SAN Volume Controller, Tivoli Storage Productivity Center, and Director-supported DS3000, DS4000 and DS5000 storage subsystems. Refer to IBM Software interoperability matrix for details.

Download the SMI-S provider and the installation instructions from the following website. First, register a free account and obtain login credentials and then download the installer.

<http://support.netapp.com/NOW/apbu/oemcp/protcd/>.

Appendix G. Accessibility

The information in this appendix describes documentation accessibility and accessibility features in Storage Manager.

Document format

The publications for this product are in Adobe Portable Document Format (PDF) and must be compliant with accessibility standards. If you experience difficulties when you use the PDF files and want to request a web-based format or accessible PDF document for a publication, direct your mail to the following address:

Information Development
IBM Corporation
205/A015
3039 E. Cornwallis Road
P.O. Box 12195
Research Triangle Park, North Carolina 27709-2195
U.S.A.

In the request, be sure to include the publication part number and title.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Accessibility features in Storage Manager

This section provides information about alternate keyboard navigation, which is a Storage Manager accessibility feature. Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

With the alternate keyboard operations that are described in this section, you can use keys or key combinations to perform Storage Manager tasks and initiate many menu actions that can also be done with a mouse.

Note: In addition to the keyboard operations that are described in this section, the Storage Manager version 9.14 - 10.10 (and later) software installation packages for Windows include a screen reader software interface.

To enable the screen reader, select **Custom Installation** when using the installation wizard to install Storage Manager 9.14 - 10.10 (or later) on a Windows host/management station. Then, in the Select Product Features window, select **Java Access Bridge**, in addition to the other required host software components.

Table 57 on page 314 defines the keyboard operations that enable you to navigate, select, or activate user interface components. The following terms are used in the table:

- *Navigate* means to move the input focus from one user interface component to another.
- *Select* means to choose one or more components, typically for a subsequent action.

- *Activate* means to carry out the action of a particular component.

Note: In general, navigation between components requires the following keys:

- **Tab** - Moves keyboard focus to the next component or to the first member of the next group of components
- **Shift-Tab** - Moves keyboard focus to the previous component or to the first component in the previous group of components
- **Arrow keys** - Move keyboard focus within the individual components of a group of components

Table 57. Storage Manager alternate keyboard operations

Short cut	Action
F1	Open the Help.
F10	Move keyboard focus to main menu bar and post first menu; use the arrow keys to navigate through the available options.
Alt+F4	Close the management window.
Alt+F6	Move keyboard focus between dialogs (non-modal) and between management windows.
Alt+ underlined letter	<p>Access menu items, buttons, and other interface components with the keys associated with the underlined letters.</p> <p>For the menu options, select the Alt + underlined letter combination to access a main menu, and then select the underlined letter to access the individual menu item.</p> <p>For other interface components, use the Alt + underlined letter combination.</p>
Ctrl+F1	Display or conceal a tool tip when keyboard focus is on the toolbar.
Spacebar	Select an item or activate a hyperlink.
Ctrl+Spacebar (Contiguous/Non-contiguous) AMW Logical/Physical View	<p>Select multiple drives in the Physical View.</p> <p>To select multiple drives, select one drive by pressing Spacebar, and then press Tab to switch focus to the next drive you want to select; press Ctrl+Spacebar to select the drive.</p> <p>If you press Spacebar alone when multiple drives are selected then all selections are removed.</p> <p>Use the Ctrl+Spacebar combination to deselect a drive when multiple drives are selected.</p> <p>This behavior is the same for contiguous and non-contiguous selection of drives.</p>
End, Page Down	Move keyboard focus to the last item in the list.
Esc	Close the current dialog. Does not require keyboard focus.
Home, Page Up	Move keyboard focus to the first item in the list.
Shift+Tab	Move keyboard focus through components in the reverse direction.

Table 57. Storage Manager alternate keyboard operations (continued)

Short cut	Action
Ctrl+Tab	Move keyboard focus from a table to the next user interface component.
Tab	Navigate keyboard focus between components or select a hyperlink.
Down arrow	Move keyboard focus down one item in the list.
Left arrow	Move keyboard focus to the left.
Right arrow	Move keyboard focus to the right.
Up arrow	Move keyboard focus up one item in the list.

Notices

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

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

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

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

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
Almaden Research
650 Harry Road
Bldg 80, D3-304, Department 277
San Jose, CA 95120-6099
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

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. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. 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 International Business Machines Corporation in the United States, other countries, or both:

IBM
AIX
eServer
FlashCopy
Netfinity
POWER
Series p
RS/6000
TotalStorage

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Intel, Intel Xeon, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

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

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.

Important notes

Processor speed indicates the internal clock speed of the microprocessor; other factors also affect application performance.

This product is not intended to be connected directly or indirectly by any means whatsoever to interfaces of public telecommunications networks.

CD or DVD drive speed is the variable read rate. Actual speeds vary and are often less than the possible maximum.

When referring to processor storage, real and virtual storage, or channel volume, KB stands for 1024 bytes, MB stands for 1,048,576 bytes, and GB stands for 1,073,741,824 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1,000,000 bytes, and GB stands for 1,000,000,000 bytes. Total user-accessible capacity can vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives that are available from IBM.

Maximum memory might require replacement of the standard memory with an optional memory module.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven[®], including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software might differ from its retail version (if available) and might not include user manuals or all program functionality.

Glossary

This glossary provides definitions for the terminology and abbreviations used in IBM System Storage publications.

If you do not find the term you are looking for, see the *IBM Glossary of Computing Terms* located at the following website:

<http://www.ibm.com/ibm/terminology>

This glossary also includes terms and definitions from:

- *Information Technology Vocabulary* by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.
- *IBM Glossary of Computing Terms*. New York: McGraw-Hill, 1994.

The following cross-reference conventions are used in this glossary:

See Refers you to (a) a term that is the expanded form of an abbreviation or acronym, or (b) a synonym or more preferred term.

See also Refers you to a related term.

Abstract Windowing Toolkit (AWT)

In Java programming, a collection of GUI components that were implemented using native-platform versions of the components. These components provide that subset of functionality which is common to all operating system environments.

accelerated graphics port (AGP)

A bus specification that gives low-cost 3D graphics cards faster access to main memory on personal computers than the usual peripheral component interconnect

(PCI) bus. AGP reduces the overall cost of creating high-end graphics subsystems with existing system memory.

access logical drive

A logical drive that allows the host-agent to communicate with the controllers in the storage subsystem.

adapter

A printed circuit assembly that transmits user data input/output (I/O) between the internal bus of the host system and the external Fibre Channel (FC) link and vice versa. Also called an I/O adapter, host adapter, or FC adapter.

advanced technology (AT) bus architecture

A bus standard for IBM compatibles. It extends the XT bus architecture to 16 bits and also allows for bus mastering, although only the first 16 MB of main memory are available for direct access.

agent A server program that receives virtual connections from the network manager (the client program) in a Simple Network Management Protocol-Transmission Control Protocol/Internet Protocol (SNMP-TCP/IP) network-managing environment.

AGP See *accelerated graphics port*.

AL_PA

See *arbitrated loop physical address*.

arbitrated loop

One of three existing Fibre Channel topologies, in which 2 - 126 ports are interconnected serially in a single loop circuit. Access to the Fibre Channel Arbitrated Loop (FC-AL) is controlled by an arbitration scheme. The FC-AL topology supports all classes of service and guarantees in-order delivery of FC frames when the originator and responder are on the same FC-AL. The default topology for the disk array is arbitrated loop. An arbitrated loop is sometimes referred to as a Stealth Mode.

arbitrated loop physical address (AL_PA)

An 8-bit value used to identify a

- participating device in an arbitrated loop. A loop can have one or more AL_PAs.
- array** A collection of Fibre Channel or SATA hard drives that are logically grouped together. All the drives in the array are assigned the same RAID level. An array is sometimes referred to as a "RAID set." See also *redundant array of independent disks (RAID)*, *RAID level*.
- asynchronous write mode**
In remote mirroring, an option that allows the primary controller to return a write I/O request completion to the host server before data has been successfully written by the secondary controller. See also *synchronous write mode*, *remote mirroring*, *Global Copy*, *Global Mirroring*.
- AT** See *advanced technology (AT) bus architecture*.
- ATA** See *AT-attached*.
- AT-attached**
Peripheral devices that are compatible with the original IBM AT computer standard in which signals on a 40-pin AT-attached (ATA) ribbon cable followed the timings and constraints of the Industry Standard Architecture (ISA) system bus on the IBM PC AT computer. Equivalent to integrated drive electronics (IDE).
- Auto Drive Transfer (ADT)**
A function that provides automatic failover in case of controller failure on a storage subsystem.
- ADT** See *Auto Drive Transfer*.
- AWT** See *Abstract Windowing Toolkit*.
- Basic Input/Output System (BIOS)**
The code that controls basic hardware operations, such as interactions with diskette drives, hard disk drives, and the keyboard.
- BIOS** See *basic input/output system*.
- BOOTP**
See *bootstrap protocol*.
- Bootstrap Protocol (BOOTP)**
A protocol that allows a client to find both its Internet Protocol (IP) address and the name of a file from a server on the network.
- bridge** A storage area network (SAN) device that provides physical and transport conversion, such as Fibre Channel to small computer system interface (SCSI) bridge.
- bridge group**
A bridge and the collection of devices connected to it.
- broadcast**
The simultaneous transmission of data to more than one destination.
- cathode ray tube (CRT)**
A display device in which controlled electron beams are used to display alphanumeric or graphical data on an electroluminescent screen.
- client** A software program or computer that requests services from a server. Multiple clients can share access to a common server.
- command**
A statement used to initiate an action or start a service. A command consists of the command name abbreviation, and its parameters and flags if applicable. A command can be issued by typing it on a command line or selecting it from a menu.
- community string**
The name of a community contained in each Simple Network Management Protocol (SNMP) message.
- concurrent download**
A method of downloading and installing firmware that does not require the user to stop I/O to the controllers during the process.
- CRC** See *cyclic redundancy check*.
- CRT** See *cathode ray tube*.
- CRU** See *customer replaceable unit*.
- customer replaceable unit (CRU)**
An assembly or part that a customer can replace. Contrast with *field replaceable unit (FRU)*.
- cyclic redundancy check (CRC)**
(1) A redundancy check in which the check key is generated by a cyclic algorithm. (2) An error detection technique performed at both the sending and receiving stations.

dac See *disk array controller*.

dar See *disk array router*.

DASD
See *direct access storage device*.

data striping
Storage process in which information is split into blocks (a fixed amount of data) and the blocks are written to (or read from) a series of disks in parallel.

default host group
A logical collection of discovered host ports, defined host computers, and defined host groups in the storage-partition topology that fulfill the following requirements:

- Are not involved in specific logical drive-to-LUN mappings
- Share access to logical drives with default logical drive-to-LUN mappings

device type
Identifier used to place devices in the physical map, such as the switch, hub, or storage.

DHCP See *Dynamic Host Configuration Protocol*.

direct access storage device (DASD)
A device in which access time is effectively independent of the location of the data. Information is entered and retrieved without reference to previously accessed data. (For example, a disk drive is a DASD, in contrast with a tape drive, which stores data as a linear sequence.) DASDs include both fixed and removable storage devices.

direct memory access (DMA)
The transfer of data between memory and an input/output (I/O) device without processor intervention.

disk array controller (dac)
The device, such as a Redundant Array of Independent Disks (RAID), that manages one or more disk arrays and provides functions. See also *disk array router*.

disk array router (dar)
A router that represents an entire array, including current and deferred paths to all logical unit numbers (LUNs) (hdisks on AIX). See also *disk array controller*.

DMA See *direct memory access*.

domain
The most significant byte in the node port (N_port) identifier for the Fibre Channel (FC) device. It is not used in the Fibre Channel-small computer system interface (FC-SCSI) hardware path ID. It is required to be the same for all SCSI targets logically connected to an FC adapter.

drive channels
The DS4200, DS4700, and DS4800 subsystems use dual-port drive channels that, from the physical point of view, are connected in the same way as two drive loops. However, from the point of view of the number of drives and enclosures, they are treated as a single drive loop instead of two different drive loops. A group of storage expansion enclosures are connected to the DS3000 or DS4000 storage subsystems using a drive channel from each controller. This pair of drive channels is referred to as a redundant drive channel pair.

drive loops
A drive loop consists of one channel from each controller combined to form one pair of redundant drive channels or a redundant drive loop. Each drive loop is associated with two ports. (There are two drive channels and four associated ports per controller.) For the DS4800, drive loops are more commonly referred to as drive channels. See *drive channels*.

DRAM
See *dynamic random access memory*.

Dynamic Host Configuration Protocol (DHCP)
A communications protocol that is used to centrally manage configuration information. For example, DHCP automatically assigns IP addresses to computers in a network.

dynamic random access memory (DRAM)
Storage in which the cells require repetitive application of control signals to retain stored data.

ECC See *error correction code*.

EEPROM
See *electrically erasable programmable read-only memory*.

EISA See *Extended Industry Standard Architecture*.

electrically erasable programmable read-only memory (EEPROM)

A type of memory chip which can retain its contents without consistent electrical power. Unlike the PROM which can be programmed only once, the EEPROM can be erased electrically. Because it can only be reprogrammed a limited number of times before it wears out, it is appropriate for storing small amounts of data that are changed infrequently.

electrostatic discharge (ESD)

The flow of current that results when objects that have a static charge come into close enough proximity to discharge.

environmental service module (ESM) canister

A component in a storage expansion enclosure that monitors the environmental condition of the components in that enclosure. Not all storage subsystems have ESM canisters.

E_port See *expansion port*.

error correction coding (ECC)

A code appended to a data block that has the capability to detect and correct multiple bit errors within the block. Most ECCs are characterized by the maximum number of errors they can detect and correct.

ESD See *electrostatic discharge*.

ESM canister

See *environmental service module canister*.

automatic ESM firmware synchronization

When you install a new ESM into an existing storage expansion enclosure in a DS3000 or DS4000 storage subsystem that supports automatic ESM firmware synchronization, the firmware in the new ESM is automatically synchronized with the firmware in the existing ESM.

EXP See *storage expansion enclosure*.

expansion port (E_port)

In the building of a larger switch fabric, a port used as an inter-switch expansion port to connect to the E_port of another switch.

Extended Industry Standard Architecture (EISA)

The PC bus standard that extends the AT bus (ISA bus) to 32 bits and provides support for bus master. It was announced in 1988 as a 32-bit alternative to the Micro

Channel that would preserve investment in existing boards. PC and AT adapters (ISA adapters) can plug into an EISA bus. See also *Industry Standard Architecture*.

fabric A Fibre Channel entity which interconnects and facilitates logins of N_ports attached to it. The fabric is responsible for routing frames between source and destination N_ports using address information in the frame header. A fabric can be as simple as a point-to-point channel between two N_ports, or as complex as a frame-routing switch that provides multiple and redundant internal pathways within the fabric between F_ports.

fabric port (F_port)

In a fabric, an access point for connecting a user's N_port. An F_port facilitates N_port logins to the fabric from nodes connected to the fabric. An F_port is addressable by the N_port connected to it. See also *fabric*.

FC See *Fibre Channel*.

FC-AL See *arbitrated loop*.

feature enable identifier

A unique identifier for the storage subsystem, which is used in the process of generating a premium feature key. See also *premium feature key*.

Fibre Channel (FC)

A technology for transmitting data between computer devices. It is especially suited for attaching computer servers to shared storage devices and for interconnecting storage controllers and drives. FC supports point-to-point, arbitrated loop, and switched topologies.

Fibre Channel Arbitrated Loop (FC-AL)

See *arbitrated loop*.

Fibre Channel Protocol (FCP) for small computer system interface (SCSI)

A high-level Fibre Channel mapping layer (FC-4) that uses lower-level Fibre Channel (FC-PH) services to transmit SCSI commands, data, and status information between a SCSI initiator and a SCSI target across the FC link with FC frame and sequence formats.

field replaceable unit (FRU)

An assembly that is replaced in its

entirety when any one of its components fails. In some cases, a field replaceable unit might contain other field replaceable units. Contrast with *customer replaceable unit (CRU)*.

FlashCopy

An optional feature of the Storage System DS family that can make an instant copy of data, that is, a point-in-time copy of a logical drive.

F_port See *fabric port*.

FRU See *field replaceable unit*.

GBIC See *gigabit interface converter*

gigabit interface converter (GBIC)

An encoding/decoding device that is a class-1 laser component assembly with transmitting and receiving receptacles that connect to fiber-optic cables. GBICs perform a serial optical-to-electrical and electrical-to-optical conversion of the signal. The GBICs in the switch can be hot-swapped. See also *small form-factor pluggable*.

Global Copy

Refers to a remote logical drive mirror pair that is set up using asynchronous write mode without the write consistency group option. This is also referred to as "Asynchronous Mirroring without Consistency Group." Global Copy does not make sure that write requests to multiple primary logical drives are carried out in the same order on the secondary logical drives as they are on the primary logical drives. If it is critical that writes to the primary logical drives are carried out in the same order in the appropriate secondary logical drives, Global Mirroring must be used instead of Global Copy. See also *asynchronous write mode*, *Global Mirroring*, *remote mirroring*, *Metro Mirroring*.

Global Mirror

An optional capability of the remote mirror and copy feature that provides a two-site extended-distance remote copy. Data that is written by the host to the storage unit at the local site is automatically maintained at the remote site. See also *asynchronous write mode*, *Global Copy*, *remote mirroring*, *Metro Mirroring*.

graphical user interface (GUI)

A type of computer interface that presents a visual metaphor of a real-world scene, often of a desktop, by combining high-resolution graphics, pointing devices, menu bars and other menus, overlapping windows, icons and the object-action relationship.

GUI See *graphical user interface*.

HBA See *host bus adapter*.

hdisk An AIX term representing a logical unit number (LUN) on an array.

heterogeneous host environment

A host system in which multiple host servers, which use different operating systems with their own unique disk storage subsystem settings, connect to the same storage subsystem at the same time. See also *host*.

host A system that is directly attached to the storage subsystem through a Fibre Channel input/output (I/O) path. This system is used to serve data (typically in the form of files) from the storage subsystem. A system can be both a management station and a host simultaneously.

host bus adapter (HBA)

An interface card that connects a host bus, such as a peripheral component interconnect (PCI) bus, to the storage area network.

host computer

See *host*.

host group

An entity in the storage partition topology that defines a logical collection of host computers that require shared access to one or more logical drives.

host port

Ports that physically reside on the host adapters and are automatically discovered by the Storage Manager software. To give a host computer access to a partition, its associated host ports must be defined.

hot-swap

Pertaining to a device that is capable of being replaced while the system is on.

hub In a network, a point at which circuits are either connected or switched. For

example, in a star network, the hub is the central node; in a star/ring network, it is the location of wiring concentrators.

IBMSAN driver

The device driver that is used in a Novell NetWare environment to provide multipath input/output (I/O) support to the storage controller.

IC See *integrated circuit*.

IDE See *integrated drive electronics*.

in-band

Transmission of management protocol over the Fibre Channel transport.

Industry Standard Architecture (ISA)

Unofficial name for the bus architecture of the IBM PC/XT personal computer. This bus design included expansion slots for plugging in various adapter boards. Early versions had an 8-bit data path, later expanded to 16 bits. The "Extended Industry Standard Architecture" (EISA) further expanded the data path to 32 bits. See also *Extended Industry Standard Architecture*.

initial program load (IPL)

The process that loads the system programs from the system auxiliary storage, checks the system hardware, and prepares the system for user operations. Also referred to as a system restart, system startup, and boot.

integrated circuit (IC)

A microelectronic semiconductor device that consists of many interconnected transistors and other components. ICs are constructed on a small rectangle cut from a silicon crystal or other semiconductor material. The small size of these circuits allows high speed, low power dissipation, and reduced manufacturing cost compared with board-level integration. Also known as a *chip*.

integrated drive electronics (IDE)

A disk drive interface based on the 16-bit IBM personal computer Industry Standard Architecture (ISA) in which the controller electronics reside on the drive itself, eliminating the need for a separate adapter card. Also known as an Advanced Technology Attachment Interface (ATA).

Internet Protocol (IP)

A protocol that routes data through a network or interconnected networks. IP acts as an intermediary between the higher protocol layers and the physical network.

Internet Protocol (IP) address

A unique address for a device or logical unit on a network that uses the IP standard. For example, 9.67.97.103 is an IP address.

interrupt request (IRQ)

An input found on a processor that causes it to suspend normal instruction execution temporarily and to start executing an interrupt handler routine.

IP See *Internet Protocol*.

IPL See *initial program load*.

IRQ See *interrupt request*.

ISA See *Industry Standard Architecture*.

Java runtime environment (JRE)

A subset of a Java developer kit that contains the core executable programs and files that constitute the standard Java platform. The JRE includes the Java virtual machine (JVM), core classes, and supporting files.

JRE See *Java Runtime Environment*.

label A discovered or user-entered property value that is displayed underneath each device in the Physical and Data Path maps.

LAN See *local area network*.

LBA See *logical block address*.

local area network (LAN)

A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

logical block address (LBA)

The address of a logical block. Logical block addresses are typically used in host I/O commands. The SCSI disk command protocol, for example, uses logical block addresses.

logical partition (LPAR)

A subset of a single system that contains resources (processors, memory, and input/output devices). A logical partition

operates as an independent system. If hardware requirements are met, multiple logical partitions can exist within a system.

A fixed-size portion of a logical drive. A logical partition is the same size as the physical partitions in its array. Unless the logical drive of which it is a part is mirrored, each logical partition corresponds to, and its contents are stored on, a single physical partition.

One to three physical partitions (copies). The number of logical partitions within a logical drive is variable.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

loop address

The unique ID of a node in Fibre Channel loop topology sometimes referred to as a loop ID.

loop group

A collection of storage area network (SAN) devices that are interconnected serially in a single loop circuit.

loop port

A port used to connect a node to a Fibre Channel Arbitrated Loop (FC-AL).

LPAR See *logical partition*.

LUN See *logical unit number*.

MAC See *medium access control*.

Management Information Base (MIB)

In the Simple Network Management Protocol (SNMP), a database of objects that can be queried or set by a network management system.

A definition for management information that specifies the information available from a host or gateway and the operations allowed.

management station

A system that is used to manage the storage subsystem. A management station does not need to be attached to the storage subsystem through the Fibre Channel input/output (I/O) path.

man page

In UNIX systems, one page of online documentation. Each UNIX command, utility, and library function has an associated man page.

MCA See *micro channel architecture*.

media scan

A media scan is a background process that runs on all logical drives in the storage subsystem for which it has been enabled, providing error detection on the drive media. The media scan process scans all logical drive data to verify that it can be accessed, and optionally scans the logical drive redundancy information.

Media Access Control (MAC)

In networking, the lower of two sublayers of the Open Systems Interconnection model data link layer. The MAC sublayer handles access to shared media, such as whether token passing or contention will be used.

metro mirror

A function of the remote mirror and copy feature that constantly updates a secondary copy of a logical drive to match changes made to a source logical drive. See also *remote mirroring*, *Global Mirroring*.

MIB See *management information base*.

Micro Channel architecture (MCA)

The rules that define how subsystems and adapters use the Micro Channel bus in a computer. MCA defines the services that each subsystem can or must provide.

Microsoft Cluster Server (MSCS)

A technology that provides high availability by grouping computers into MSCS clusters. If one of the computers in the cluster hits any one of a range of problems, MSCS shuts down the disrupted application in an orderly manner, transfers its state data to another computer in the cluster, and re-initiates the application there.

mini hub

An interface card or port device that receives short-wave fiber channel GBICs or SFPs. These devices enable redundant Fibre Channel connections from the host computers, either directly or through a Fibre Channel switch or managed hub,

over optical fiber cables to the DS3000 and DS4000 Storage Server controllers. Each DS3000 and DS4000 controller is responsible for two mini hubs. Each mini hub has two ports. Four host ports (two on each controller) provide a cluster solution without use of a switch. Two host-side mini hubs are shipped as standard. See also *host port*, *gigabit interface converter (GBIC)*, *small form-factor pluggable (SFP)*.

mirroring

A fault-tolerance technique in which information on a hard disk is duplicated on additional hard disks. See also *remote mirroring*.

model The model identification that is assigned to a device by its manufacturer.

MSCS See *Microsoft Cluster Server*.

network management station (NMS)

In the Simple Network Management Protocol (SNMP), a station that runs management application programs that monitor and control network elements.

NMI See *non-maskable interrupt*.

NMS See *network management station*.

non-maskable interrupt (NMI)

A hardware interrupt that another service request cannot overrule (mask). An NMI bypasses and takes priority over interrupt requests generated by software, the keyboard, and other such devices and is issued to the microprocessor only in disastrous circumstances, such as severe memory errors or impending power failures.

node A physical device that allows for the transmission of data within a network.

node port (N_port)

A Fibre Channel defined hardware entity that performs data communications over the Fibre Channel link. It is identifiable by a unique worldwide name. It can act as an originator or a responder.

nonvolatile storage (NVS)

A storage device whose contents are not lost when power is cut off.

N_port

See *node port*.

NVS See *nonvolatile storage*.

NVSRAM

Nonvolatile storage random access memory. See *nonvolatile storage*.

Object Data Manager (ODM)

An AIX proprietary storage mechanism for ASCII stanza files that are edited as part of configuring a drive into the kernel.

ODM See *Object Data Manager*.

out-of-band

Transmission of management protocols outside of the Fibre Channel network, typically over Ethernet.

partitioning

See *storage partition*.

parity check

A test to determine whether the number of ones (or zeros) in an array of binary digits is odd or even.

A mathematical operation on the numerical representation of the information communicated between two pieces. For example, if parity is odd, any character represented by an even number has a bit added to it, making it odd, and an information receiver checks that each unit of information has an odd value.

PCI local bus

See *peripheral component interconnect local bus*.

PDF See *portable document format*.

performance event

Event related to thresholds set on storage area network (SAN) performance.

Peripheral Component Interconnect local bus (PCI local bus)

A local bus for PCs, from Intel, that provides a high-speed data path between the CPU and up to 10 peripherals (video, disk, network, and so on). The PCI bus coexists in the PC with the Industry Standard Architecture (ISA) or Extended Industry Standard Architecture (EISA) bus. ISA and EISA boards plug into an IA or EISA slot, while high-speed PCI controllers plug into a PCI slot. See also *Industry Standard Architecture*, *Extended Industry Standard Architecture*.

polling delay

The time in seconds between successive discovery processes during which discovery is inactive.

port

A part of the system unit or remote controller to which cables for external devices (such as display stations, terminals, printers, switches, or external storage units) are attached. The port is an access point for data entry or exit. A device can contain one or more ports.

portable document format (PDF)

A standard specified by Adobe Systems, Incorporated, for the electronic distribution of documents. PDF files are compact; can be distributed globally via e-mail, the Web, intranets, CD-ROM or DVD-ROM; and can be viewed with the Acrobat Reader.

premium feature key

A file that the storage subsystem controller uses to enable an authorized premium feature. The file contains the feature enable identifier of the storage subsystem for which the premium feature is authorized, and data about the premium feature. See also *feature enable identifier*.

private loop

A Fibre Channel Arbitrated Loop (FC-AL) with no fabric attachment. See also *arbitrated loop*.

program temporary fix (PTF)

For System i, System p, and System z products, a fix that is tested by IBM and is made available to all customers.

PTF See *program temporary fix*.

RAID See *redundant array of independent disks (RAID)*.

RAID level

An array RAID level is a number that refers to the method used to achieve redundancy and fault tolerance in the array. See also *array*, *redundant array of independent disks (RAID)*.

RAID set

See *array*.

RAM See *random access memory*.

random access memory (RAM)

Computer memory in which any storage location can be accessed directly. Contrast with *DASD*.

RDAC

See *redundant disk array controller*.

read-only memory (ROM)

Memory in which stored data cannot be changed by the user except under special conditions.

recoverable virtual shared disk (RVSD)

A virtual shared disk on a server node configured to provide continuous access to data and file systems in a cluster.

Redundant Array of Independent Disks (RAID)

A collection of two or more physical disk drives (or an *array*) that present to the host an image of one or more logical disk drives. In the event of a physical device failure, the data can be read or regenerated from the other disk drives in the array due to data redundancy. See also *array*, *parity check*, *mirroring*, *RAID level*, *striping*.

redundant disk array controller (RDAC)

In hardware, a redundant set of controllers (either active/passive or active/active).

In software, a layer that manages the input/output (I/O) through the active controller during normal operation and transparently reroutes I/Os to the other controller in the redundant set if a controller or I/O path fails.

remote mirroring

Online, real-time replication of data between storage subsystems that are maintained on separate media. The Enhanced Remote Mirror Option is a premium feature that provides support for remote mirroring. See also *Global Mirroring*, *Metro Mirroring*.

ROM See *read-only memory*.

router

A computer that determines the path of network traffic flow. The path selection is made from several paths based on information obtained from specific protocols, algorithms that attempt to identify the shortest or best path, and other criteria such as metrics or protocol-specific destination addresses.

RVSD See *recoverable virtual shared disk*.

SAI See *Storage Subsystem Identifier*.

SA Identifier
See *storage subsystem Identifier*.

SAN See *storage area network*.

SATA See *serial ATA*.

scope Defines a group of controllers by their Internet Protocol (IP) addresses. A scope must be created and defined so that dynamic IP addresses can be assigned to controllers on the network.

SCSI See *small computer system interface*.

segmented loop port (SL_port)
A port that allows division of a Fibre Channel private loop into multiple segments. Each segment can pass frames around as an independent loop and can connect through the fabric to other segments of the same loop.

sense data
Data sent with a negative response, indicating the reason for the response.

Data describing an I/O error. Sense data is presented to a host system in response to a sense request command.

serial ATA
The standard for a high-speed alternative to small computer system interface (SCSI) hard drives. The SATA-1 standard is equivalent in performance to a 10 000 RPM SCSI drive.

Serial Storage Architecture (SSA)
An American National Standards Institute (ANSI) standard, implemented by IBM, for a high-speed serial interface that provides point-to-point connection for peripherals, such as storage subsystems. SSA, which is compatible with small computer system interface (SCSI) devices, allows full-duplex packet multiplexed serial data transfers at rates of 20 Mbps in each direction.

server A software program or a computer that provides services to other software programs or other computers.

server/device events
Events that occur on the server or a designated device that meet criteria that the user sets.

SFP See *small form-factor pluggable*.

Simple Network Management Protocol (SNMP)
A set of protocols for monitoring systems and devices in complex networks. Information about managed devices is defined and stored in a Management Information Base (MIB).

SL_port
See *segmented loop port*.

SMagent
The Storage Manager optional Java-based host-agent software, which can be used on Microsoft Windows, Novell NetWare, AIX, HP-UX, Solaris, and Linux on POWER host systems to manage storage subsystems through the host Fibre Channel connection.

SMclient
The Storage Manager client software, which is a Java-based graphical user interface (GUI) that is used to configure, manage, and troubleshoot storage servers and storage expansion enclosures in a storage subsystem. SMclient can be used on a host system or on a management station.

SMruntime
A Java compiler for the SMclient.

SMutil
The Storage Manager utility software that is used on Microsoft Windows, AIX, HP-UX, Solaris, and Linux on POWER host systems to register and map new logical drives to the operating system. In Microsoft Windows, it also contains a utility to flush the cached data of the operating system for a particular drive before creating a FlashCopy.

Small Computer System Interface (SCSI)
An ANSI-standard electronic interface that allows personal computers to communicate with peripheral hardware, such as disk drives, tape drives, CD-ROM or DVD-ROM drives, printers, and scanners faster and more flexibly than previous interfaces.

A standard hardware interface that enables a variety of peripheral devices to communicate with one another.

small form-factor pluggable (SFP)
An optical transceiver that is used to

- convert signals between optical fiber cables and switches. An SFP is smaller than a gigabit interface converter (GBIC). See also *gigabit interface converter*.
- SNMP**
See *Simple Network Management Protocol* and *SNMPv1*.
- SNMP trap event**
An event notification sent by the SNMP agent that identifies conditions, such as thresholds, that exceed a predetermined value. See also *Simple Network Management Protocol*.
- SNMPv1**
The original standard for SNMP is now referred to as SNMPv1, as opposed to SNMPv2, a revision of SNMP. See also *Simple Network Management Protocol*.
- SRAM**
See *static random access memory*.
- SSA** See *serial storage architecture*.
- static random access memory (SRAM)**
Random access memory based on the logic circuit known as flip-flop. It is called static because it retains a value as long as power is supplied, unlike dynamic random access memory (DRAM), which must be regularly refreshed. It is however, still volatile, meaning that it can lose its contents when the power is turned off.
- storage area network (SAN)**
A dedicated storage network tailored to a specific environment, combining servers, storage products, networking products, software, and services. See also *fabric*.
- Storage Subsystem Identifier (SAI or SA Identifier)**
The Storage Subsystem Identifier is the identification value used by the Storage Manager host software (SMClient) to uniquely identify each managed storage server. The Storage Manager SMClient program maintains Storage Subsystem Identifier records of previously-discovered storage servers in the host resident file, which allows it to retain discovery information in a persistent fashion.
- storage expansion enclosure (EXP), or storage enclosure**
A feature that can be connected to a system unit to provide additional storage and processing capacity.
- storage partition**
Storage subsystem logical drives that are visible to a host computer or are shared among host computers that are part of a host group.
- storage partition topology**
In the Storage Manager client, the Topology view of the Mappings window displays the default host group, the defined host group, the host computer, and host-port nodes. The host port, host computer, and host group topological elements must be defined to grant access to host computers and host groups using logical drive-to-LUN mappings.
- striping**
See *data striping*.
- subnet**
A network divided into smaller independent subgroups, which still are interconnected.
- sweep method**
A method of sending Simple Network Management Protocol (SNMP) requests for information to all the devices on a subnet by sending the request to every device in the network.
- switch** A Fibre Channel device that provides full bandwidth per port and high-speed routing of data with link-level addressing.
- switch group**
A switch and the collection of devices connected to it that are not in other groups.
- switch zoning**
See *zoning*.
- synchronous write mode**
In remote mirroring, an option that requires the primary controller to wait for the acknowledgment of a write operation from the secondary controller before returning a write I/O request completion to the host. See also *asynchronous write mode*, *remote mirroring*, *Metro Mirroring*.
- system name**
Device name assigned by the vendor third-party software.
- TCP** See *Transmission Control Protocol*.

TCP/IP

See *Transmission Control Protocol/Internet Protocol*.

terminate and stay resident program (TSR program)

A program that installs part of itself as an extension of DOS when it is executed.

topology

The physical or logical mapping of the location of networking components or nodes within a network. Common network topologies include bus, ring, star, and tree. The three Fibre Channel topologies are fabric, arbitrated loop, and point-to-point. The default topology for the disk array is arbitrated loop.

TL_port

See *translated loop port*.

transceiver

In communications, the device that connects the transceiver cable to the Ethernet coaxial cable. The transceiver is used to transmit and receive data. Transceiver is an abbreviation of transmitter-receiver.

translated loop port (TL_port)

A port that connects to a private loop and allows connectivity between the private loop devices and off loop devices (devices not connected to that particular TL_port).

Transmission Control Protocol (TCP)

A communication protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for internetwork protocol. TCP provides a reliable host-to-host protocol in packet-switched communication networks and in interconnected systems of such networks.

Transmission Control Protocol/Internet Protocol (TCP/IP)

A set of communication protocols that provide peer-to-peer connectivity functions for both local and wide-area networks.

trap

In the Simple Network Management Protocol (SNMP), a message sent by a managed node (agent function) to a management station to report an exception condition.

trap recipient

Receiver of a forwarded Simple Network Management Protocol (SNMP) trap. Specifically, a trap receiver is defined by an Internet Protocol (IP) address and port to which traps are sent. Presumably, the actual recipient is a software application running at the IP address and listening to the port.

TSR program

See *terminate and stay resident program*.

uninterruptible power supply

A source of power from a battery installed between the commercial power and the system that keeps the system running, if a commercial power failure occurs, until it can complete an orderly end to system processing.

user action events

Actions that the user takes, such as changes in the storage area network (SAN), changed settings, and so on.

worldwide port name (WWPN)

A unique 64-bit identifier associated with a switch. The WWPN is assigned in an implementation-independent and protocol-independent manner.

worldwide name (WWN)

A 64-bit, unsigned, unique name identifier that is assigned to each Fibre Channel port.

WORM

See *write-once read-many*.

Write Once Read Many (WORM)

Any type of storage medium to which data can be written only a single time, but can be read from any number of times. After the data is recorded, it cannot be altered.

WWN See *worldwide name*.**zoning**

In Fibre Channel environments, the grouping of multiple ports to form a virtual, private, storage network. Ports that are members of a zone can communicate with each other, but are isolated from ports in other zones.

A function that allows segmentation of nodes by address, name, or physical port and is provided by fabric switches or hubs.

Index

A

- access volumes 6, 303
- accessibility
 - document 313
 - Storage Manager features 313
- adapter
 - See* HBAs
- ADT 170, 171
- AIX 129
 - cluster services 291
 - error log 266
 - FCP disk array errors 266
 - HBA hot swap, completing the 179
 - HBA hot swap, preparing for 172
 - hot-swap HBA, replacing 174
 - logical drives, redistributing 169
 - mapping new WWPNs to storage subsystems 179
 - Object Data Manager (ODM)
 - attributes
 - dac devices 301
 - dar devices 301
 - definitions 301
 - hdisk devices 301
 - lsattr command 306
 - viewing and setting 301
 - AIX multipath drivers 129
- alert notifications, setting 37
- arrays 74
 - creating 67
 - definition 67
- attributes
 - dac 306
 - dar 306
 - hdisk 164, 306
 - LUN 164, 306
- automatic ESM firmware synchronization
 - defined 48
 - Event Monitor requirement 48
- automatic host discovery 34
- automatic storage subsystem
 - discovery 34
- autorecovery
 - verifying disabled state before HBA hot swap on AIX 172

B

- background media scan 91
- BIOS
 - settings 271

C

- cache hit
 - optimizing 98
 - percentage 98
- cache mirroring 268, 293
- cache mirroring, disabling 165

- cache read-ahead, choosing a multiplier 98
- CHAP 40
- cluster services
 - AIX requirements 293
 - HACMP ES and ESCRM 291
- cluster services, high-availability
 - AIX 291
 - AIX requirements 292
 - HP-UX requirements 299
 - MC/Service Guard 299
 - PSSP with GPFS 293
 - Solaris 300
 - Solaris requirements 300
 - system dependencies 291
- clustering
 - VMware ESX Server
 - configuration 286
- command-line interface (CLI) 100
- comments, Script Editor 100
- components, Storage Manager
 - software 2, 32
- concurrent firmware download 43, 49
- configuration 5, 212, 220
 - device driver, Linux DM-Multipath driver 129
 - devices 162
 - direct-attached 2, 6
 - DS TKLM Proxy Code server, starting, stopping, and restarting 211
 - DS TKLM Proxy Code, for external security key management 210
 - FDE drives 216
 - GPFS, PSSP, and HACMP
 - clusters 293
 - HBAs 271
 - hosts 103
 - hot-spare drives 80
 - IBM i 85
 - iSCSI host ports 40
 - iSCSI settings 39
 - MTU 42
 - network 2, 3
 - network example 4
 - network settings for iSCSI host
 - attachment 42
 - recovery 67
 - SAN-attached 2, 6
 - storage subsystem passwords 36
 - storage subsystems 6
 - types 2
- configuration types
 - Storage Manager installation 2
- controller cache memory 89
- controller firmware
 - downloading 43, 45
 - firmware
 - downloading 45
- Controller Firmware Upgrade Tool
 - device health, checking 47
 - firmware, downloading 48

- Controller Firmware Upgrade Tool
 - (*continued*)
 - log file, viewing 48
 - opening 47
 - overview 46
 - storage subsystems, adding 48
 - using 47
- controllers
 - addresses 7
 - dar 160
 - disk array 160
 - IP addresses 7
 - transfer rate, optimizing 97
- copy services 50
- cross connections
 - VMware ESX Server 289

D

- dac
 - and RDAC 160
 - attributes 306
- dar
 - and RDAC 160
 - attributes 306
- data
 - collecting before HBA hot swap on AIX 172
 - files, defragmenting 100
 - optimal segment size, choosing 99
 - redundancy 74
 - securing with FDE 192
- DCE 166
- DDC
 - See* Diagnostic Data Capture
- default host types, defining and verifying 81
- device drivers
 - description 115
 - DMP, installing 156
 - failover 115
 - HBAs 124, 125
 - Linux DM-Multipath driver 129
 - multipath 115
 - installing 120
 - RDAC 115, 148
 - RDAC failover driver on Solaris 154
 - SCSIport Miniport 125
 - Storport Miniport 124
 - Veritas DMP DSM 129
 - with HACMP cluster 293
- Device Specific Module
 - See* DSM
- devices
 - adding 36
 - configuring 162
 - identification 161
 - identifying 159
 - setting alert notifications 37
- Devices tab
 - See* Enterprise Management window

- DHCP, using 41
- DHCP/BOOTP server 8
- Diagnostic Data Capture
 - MEL events 265
 - Recovery Guru 263, 265
 - recovery steps 264
 - Script Editor 263
- direct-attached configuration
 - setting IP addresses 7
- direct-attached configurations
 - setting up 6
- discovery, automatic storage
 - subsystem 34
- disk access, minimizing 100
- disk array controller
 - See* dac
- disk array router
 - See* dar
- disk drives
 - FDE 192
 - FDE hot-spare 238
 - FDE, configuring 216
 - FDE, erasing 235
 - FDE, installing 217
 - FDE, migrating 232
 - FDE, secure erase 206
 - FDE, unlocking (external) 205
 - FDE, unlocking (local and external) 229
 - FDE, unlocking (local) 204
 - hot spares, assigning 80
 - hot spares, configuring 80
 - hot spares, restoring data from 80
- disk pool
 - creating 72
- DMP 148
 - planning for installation 156
 - preparing for installation 156
- DMP drivers 143
- DMP DSM driver 129
- documentation
 - about xi
 - accessibility 313
 - FDE best practices 242
 - notices xvii
 - related documentation resources xiii
 - statements xvii
 - Storage Manager xiii
 - Sun Solaris 148
 - Sun StorEdge 154
 - Symantec 129, 156
 - using xvii
 - Veritas 129, 156
 - VMware 289
 - websites xiii, xv
- drive firmware
 - downloading 49
 - levels, determining 44, 45
- drivers
 - rpaphp 175
- drives
 - See* disk drives
- DS TKLM Proxy Code server,
 - restarting 211
- DS TKLM Proxy Code server, supported
 - operating systems 210

- DS TKLM Proxy Code, configuring for
 - external security key management 220
- DS TKLM Proxy Code, for external
 - security key management 215
- DSM 121
- DVE 166
- Dynamic Capacity Expansion
 - See* DCE
- Dynamic Logical Drive Expansion
 - See* DVE
- Dynamic Multipathing (DMP)
 - See also* DMP
 - description 148

E

- Enhanced Global Mirroring
 - Using 88
- Enterprise Management window
 - alert notifications 37
 - Devices tab 16
 - elements 13
 - online help xiv
 - Setup tab 19
 - table view description 16
 - tree view description 16
- Enterprise Management Window
 - adding devices 36
- errors
 - FCP disk array 266
 - errors, media scan 93
- ESM firmware
 - automatic ESM firmware
 - download 48
 - automatic ESM firmware
 - synchronization 48
 - downloading 43, 48
 - levels, determining 44, 45
- Ethernet
 - Solaris requirements, cluster
 - services 300
- Ethernet MAC addresses
 - See* MAC addresses
- events
 - DDC MEL 265
- events, critical
 - descriptions of 245
 - numbers 245
 - required actions 245
 - solving problems 245
- external security key management 194,
 - 197, 205, 210, 212, 215, 220
 - configuring 220
 - DS TKLM Proxy Code server 215

F

- fabric switch environments 119
- failover driver
 - description 115
 - MPxIO 149
- Failover modes
 - Asymmetric Logical Unit Access 105
 - Automatic Volume Transfer 105
 - RDAC failover 105

- failure support
 - cluster services 291
 - DMP driver 148
 - MPxIO 148
 - RDAC driver 148
 - redistributing logical drives 169, 170, 171
- Fast!UTIL 271
- FC/SATA Intermix premium feature 50
- FCP disk array errors 266
- FDE 89, 191
 - arrays, securing 240
 - backup and recovery 242
 - best practices 242
 - boot support 242
 - disk drives 192
 - disk drives, configuring 216
 - disk drives, erasing 235
 - disk drives, installing 217
 - disk drives, migrating 232
 - disk drives, unlocking (local and external) 229
 - enabling 217
 - external security key
 - management 193, 241
 - frequently asked questions 239
 - hot-spare disk drives 238
 - hot-spare drives 242
 - key management method, choosing
 - a 193
 - local security key management 193, 241
 - log files 239
 - RAID arrays, securing 224
 - secure drives, unlocking
 - (external) 205
 - secure drives, unlocking (local) 204
 - secure erase 240
 - secure erase, using 206
 - securing data against a breach 192
 - security authorizations 207
 - security key identifier 198
 - security key management, FDE 193
 - security keys
 - creating 194
 - obtaining 194
 - using 194
 - security keys, changing
 - (external) 197
 - security keys, changing (local) 197
 - security keys, creating 194
 - security keys, obtaining 194
 - states, locked and unlocked 242
 - terminology 209
 - understanding 192
 - using with other premium
 - features 241
- feature enable identifier 52
- feature key file 53
- features
 - Fast!UTIL 271
- features, premium
 - See* premium features
- Fibre Channel
 - HBAs in a switch environment 119
- Fibre Channel I/O
 - cache-hit percentage 98

- Fibre Channel I/O (*continued*)
 - load balancing 97
- Fibre Channel switch zoning 119
- files, defragmenting 100
- firmware
 - controller xiv
 - downloading 43, 48
 - downloading with I/O 49
 - levels, determining 44, 45
 - obtaining xiv
 - supported by Storage Manager 2
 - versions 2
- FlashCopy 87
 - disk array error messages (AIX) 268
 - Enhanced 86
- Frequently asked questions, other 243
- full disk encryption
 - See* FDE

G

- General Parallel File System (GPFS) 293
- glossary 321
- GPFS 293

H

- HACMP 291
 - using 293
- hardware
 - Ethernet address 5
 - service and support xvii
 - VMware ESX Server requirements 286
- hardware initiators, iSCSI 41
- Hardware tab
 - See* Subsystem Management window
- HBAs
 - advanced settings 273
 - connecting an a Fibre Channel switch environment 119
 - default settings 272
 - device drivers 124, 125
 - Fibre Channel switch environment 119
 - hot swap, completing 179
 - hot-swap on Linux, preparing for 175
 - hot-swap, replacing 171
 - hot-swap, replacing for AIX and Linux 174
 - hot-swap, replacing on AIX 171, 172
 - in a direct-attached configuration 6
 - in a SAN-attached configuration 6
 - JNI settings 279
 - JNI settings on Solaris 279
 - on Linux, replacing 175
 - overview 118
 - PCI hotplug, replacing 178
 - QLogic settings 274, 283
 - settings 271
 - using 118
- hdisk
 - attributes 164, 306
 - queue depth, setting 164

- hdisks
 - verification 161
- head-swap, FDE storage subsystems 232
- help
 - obtaining xvi, xvii
 - websites xv
- help, online xiv
- heterogeneous environment 83
- High Availability Cluster Multi-Processing
 - See* HACMP
- host bus adapters
 - See* HBAs
- HBAs
 - setting host ports 55
- Solaris
 - QLogic settings 284
- host groups
 - defining 55, 83
- Host Mappings tab
 - See* Subsystem Management window
- host ports
 - defining 84
 - definition 83
- host ports, iSCSI 40
- host types
 - defining default 81
 - verifying 81
- host-agent software
 - stopping and restarting 163
- host-agent-managed configuration 5
- host-agent-managed, setting up 5
- host-agent-management method
 - UTM device 161
- hosts
 - AIX, devices on 160
 - automatic discovery 34
 - configuring 103
 - defining 84
 - heterogeneous 83
 - IBM i 85, 86
 - iSCSI 42
 - manual discovery 36
 - pre-installation tasks 6
 - VMware ESX Server 286
- hot_add utility 162
- hot-spare
 - FDE disk drives 238
- hot-spare drives 80
- hot-swap HBAs
 - See* HBAs, hot-swap
- HP-UX
 - cluster services, high-availability requirements 299
 - logical drives, redistributing 170
 - multipath I/O with PV-links 143, 144, 145
 - native multipathing 147
 - PV-links 143, 144, 145

I

- I/O
 - access pattern 98
 - request rate optimizing 98
 - size 98
 - write-caching 98

- I/O access pattern and I/O size 98
- I/O activity, monitoring 115
- I/O data field 96, 97
- I/O request rate
 - optimizing 98
- I/O transfer rate, optimizing 97
- IBM i
 - configuring 85
 - defining as the host type 86
 - port identifier 85
- IBM Power systems 85
- IBM Support Line xvii
- IBM System Storage Controller Firmware Upgrade Tool
 - See* Controller Firmware Upgrade Tool
- IBM System Storage Productivity Center xv
- IBM Tivoli Key Lifecycle Manager 193, 220
 - DS TKLM Proxy Code server configuration file, modifying 212
 - DS TKLM Proxy Code, configuring 220
 - DS TKLM Proxy Code, installing 215
 - for external security key management, configuring 210
- important notices 320
- in-band configuration
 - See* host-agent-managed configuration
- installation 215
 - completion procedures 34
 - configuration types 2
 - console window 31
 - DS TKLM Proxy Code on Windows, for external security key management 215
 - FDE drives 217
 - multipath drivers 120
 - network configuration 3
 - preparing for 1
 - proxy on AIX or Linux, for external security key management 215
 - sequence of 32
 - Storage Manager 27
 - Storage Manager, automatic 28
 - Storage Manager, manual 32
 - Support Monitor 27
 - VMware ESX Server configuration 285
- introduction
 - Storage Manager 1
- IP address
 - IPv6 42
- IP addresses 7
- IPv6 42
- iSCSI
 - host ports 39
 - host ports, configuring 40
 - iSNS server, using 41
 - mutual authentication permissions, entering 40
 - network settings 42
 - session, viewing or ending 40
 - settings, managing 39
 - software initiator considerations, Microsoft 43
 - statistics, viewing 41

- iSCSI (*continued*)
 - supported hardware initiators, using 41
 - target authentication, changing 40
 - target discovery, changing 40
 - target identification, changing 40
- iSNS server, using 41

J

- JNI
 - HBA settings 279
 - HBA settings on Solaris 279

K

- keys, security (FDE)
 - See* FDE

L

- least path weight policy 96
- least queue depth policy 96
- Linux
 - DCE 166
 - DVE 166
 - HBA hot swap, completing the 179
 - HBAs, preparing for hot swap 175
 - hot-swap HBA, replacing 174
 - mapping new WWPNS to storage subsystems 179
 - replacing HBAs 175
 - RHEL 5.3, with Veritas Storage Foundation 5.0 167
 - SUSE, with Veritas Storage Foundation 167
- Linux DM-Multipath driver 129
- Linux MPP drivers 138
- load balancing 301
- load_balancing attribute 301, 302
- local security key management 194, 197, 204
- LockKeyID, FDE 198
- log files 245
 - major events log 239
 - security changes 239
- logical drives 6, 303
 - configuration 77
 - creating 67, 77, 78
 - creating from free or unconfigured capacity 74
 - definition 67
 - expected usage 78
 - identifying 159
 - modification priority setting 99
 - redistributing 169, 170, 171
- lsslot tool 176
- LUNs
 - adding to an existing partition 84, 85
 - attributes 164, 306
 - checking size 169
 - mapping to a new partition 84
 - mapping to a partition on VMware ESX Server 289

M

- MAC addresses
 - identifying 8
- MAC OS 143
- management stations
 - compatible configuration types 2
 - description 1, 4
 - VMware ESX Server 286
- manual discovery 36
- Maximum Transmission Unit
 - See* MTU
- MC/Service Guard 299
- Media Scan 91
 - changing settings 91
 - duration 95
 - errors reported 93
 - overview 91
 - performance impact 93
 - settings 94
- medical imaging applications 75
- Medium Access Control (MAC) address
 - See* MAC addresses
- MEL
 - security changes 239
- Microsoft iSCSI Software Initiator 43
- Microsoft Windows MPIO 121
- Microsoft Windows MPIO/DSM 121
- Minihubs 6
- modifying the proxy configuration file for external security key management 212
- MPIO 162
- MPP drivers 138
- MPxIO 148
 - device names, changing 149
 - devices, verifying 149
 - driver, disabling 154
 - failover driver, enabling 149
 - failover path, configuring 149
 - failover path, unconfiguring 149
 - latest driver version, acquiring 149
- MTU
 - settings 42
- multi-user environments 76
- multimedia applications 75
- multipath driver 143
 - description 115
- multipath drivers 105, 129, 138, 159
 - failover 105
 - installing 120
- multipathing 43, 120, 148
 - DMP, installing on Solaris 156
 - MPxIO, using with Solaris 148
 - native, on HP-UX 147
 - PV-links, using on HP-UX 143, 144, 145
 - RDAC failover driver on Solaris 154
 - redistributing logical drives on AIX 169
 - redistributing logical drives on HP-UX 170
 - redistributing logical drives on Solaris 171
- Multiplexed I/O (MPxIO)
 - See* MPxIO
- mutual authentication permissions, entering for iSCSI 40
- My Support xviii

N

- names, storage subsystem 37
- network installation, preparing 3
- network-managed configuration 5
- network-managed, setting up 5
- networks
 - configuration example 4
 - general configuration 3
 - iSCSI settings 42
- notes, important 320
- notices xvii
 - general 317
- notifications
 - to alphanumeric pagers 37
 - to email 37
 - using SNMP traps 37
- NVSRAM firmware
 - downloading 43, 45

O

- Object Data Manager (ODM) attributes
 - definitions 301
 - initial device identification 161
- operating system
 - requirements 28
- operating systems
 - booting with SAN boot 103
 - DS TKLM Proxy Code 210
 - Solaris 149
 - supported for Storage Manager 1
- operations progress
 - viewing 79
- Other frequently asked questions 243
- out-of-band configuration
 - See* network-managed configuration

P

- packages, Storage Manager software 2, 32
- Parallel System Support Programs (PSSP) 293
- parity 74
- partitioning 55
- passwords, setting 36
- PCI core 175
- PCI Hotplug 176
- PCI hotplug HBAs 178
- PCI Hotplug tools 175
- PCI slot information 176
- performance
 - ODM attribute settings and 164
- Performance Monitor 95
- Performance Read Cache
 - Using 88
- Persistent Reservations 91
- policies, load-balancing
 - least path weight policy 96
 - least queue depth policy 96
 - round robin policy 96
- premium feature 88
- premium features 50
 - configuring 86
 - descriptions of 50
 - disabling 54

- premium features (*continued*)
 - enabling 53
 - FDE 191
 - FDE and FlashCopy 241
 - FDE and VolumeCopy 241
 - FDE, enabling 217
 - feature enable identifier 52
 - feature key file 53
 - FlashCopy 87
 - Full Disk Encryption
 - See* FDE
 - key 88, 89
 - Remote Mirror Option 88
 - storage partitioning 55, 83
 - using 86
 - VolumeCopy 88
- prerequisites
 - HP-UX
 - cluster services, high-availability 299
- prerequisites, Storage Manager client software 32
- problem solving, critical events 245
- problems, solving 245
- products, developed 317
- profile, storage subsystem 54
- proxy, installing on AIX or Linux 215
- proxy, installing on Windows 215
- proxy, uninstalling on Windows 215
- PSP 293
- PV-links
 - See* HP-UX

Q

- QLogic
 - HBA settings 271, 274, 279, 283
 - settings 284
- QLogic SANsurfer xiv
- queue depth
 - changing for AIX 165
 - changing for Windows 165
 - maximum, calculating 164
- queue depth, setting 164

R

- RAID
 - application behavior, by level 99
 - choosing levels 99
 - data redundancy 74
 - levels 74
 - securing arrays with FDE 224
- RAID level
 - application behavior 76
 - choosing 76
 - configurations 75
- RAID-0
 - described 75
 - drive failure consequences 75
- RAID-1
 - described 75
 - drive failure consequences 75
- RAID-3
 - described 75
 - drive failure consequences 75

- RAID-5
 - described 76
 - drive failure consequences 76
- RAID-6
 - dual distributed parity 76
- RDAC
 - IDs, adding 300
 - IDs, cluster services 300
- RDAC driver
 - description 115
- RDAC drivers
 - configuration files, modifying 154
 - description 148
 - Solaris, installing on 154
- readme
 - obtaining files xiv
- Recovery Guru
 - Diagnostic Data Capture 263
- remote boot
 - See* SAN boot
- Remote Mirror Option 88
- requirements
 - HP-UX
 - cluster services, high-availability 299
 - operating system 28
 - Storage Manager client software 32
- round robin policy 96

S

- SAN boot
 - configuring hosts 103
 - requirements 103
- SAN-attached configuration
 - setting up 6
- Script Editor
 - Diagnostic Data Capture 263
 - using 100
 - window 100
- SCSIport Miniport 125, 129
- secure erase, FDE 206
- security authorizations, FDE 207
- security keys
 - changing (external) 197
 - changing (local) 197
 - creating 194
 - identifier 198
 - using to unlock FDE drives 229
- security keys, FDE
 - See* FDE
- service
 - requesting xvi
- services offered in the U.S.A. 317
- session, iSCSI 40
- settings
 - advanced HBA 273
 - HBA 271
 - HBA default 272
 - media scan 94
 - modification priority 99
 - MTU 42
- Setup tab
 - Enterprise Management window 19
 - Subsystem Management window 26
- SMagent
 - software installation sequence 32

- SMclient
 - software installation sequence 32
- SMdevices utility
 - on Unix-type operating systems 160
 - on Windows 159
 - using 159
- SMrepassist utility 163
- SMruntime
 - software installation sequence 32
- SMutil
 - software installation sequence 32
- SNMP traps 37
- software
 - key license management 193
 - multipath drivers 115
 - RDAC package 148
 - service and support xvii
 - setting up controller addresses 7
 - Storage Manager components 2, 32
 - VMware ESX Server
 - requirements 286
- software versions, multiple
 - See* Subsystem Management window
- Solaris
 - cluster services requirements 300
 - cluster services, high-availability 300
 - logical drives, redistributing 171
 - Veritas 156
- SSPC (System Storage Productivity Center) xv
- statements xvii
- statistics, iSCSI 41
- Storage and Copy Services tab
 - See* Subsystem Management window
- storage area network (SAN)
 - configuration 6
- Storage Manager 89, 91
 - accessibility 313
 - command-line interface 100
 - Controller Firmware Upgrade Tool
 - using the tool 46
 - description 1
 - Enterprise Management window 13
 - event log 245
 - installation 34
 - installation sequence 32
 - installation wizard 28
 - installing for use with IBM System Storage Productivity Center xv
 - interface elements 13
 - interface elements, Storage Manager 13
 - manual installation 32
 - obtaining software xiv
 - other features 89
 - premium features 50, 86
 - problems, solving 245
 - Script Editor 100
 - setting up controller addresses 7
 - software components 2, 32
 - Subsystem Management window 19
 - supported operating systems 1
 - Task Assistant 56
 - troubleshooting 245
 - uninstalling 33
 - uninstalling on Linux, AIX, or Solaris 34

- Storage Manager 89, 91 (*continued*)
 - uninstalling on Windows 33
 - version 10.5x drive firmware
 - download 49
 - versions 2
- storage partitioning 50, 83
 - and host groups 55
- storage subsystem
 - adding 163
 - firmware levels, determining 44, 45
 - profile, saving 54
 - VMware ESX Server
 - configuration 285
- storage subsystem, for external security
 - key management 220
- storage subsystems
 - cluster services 291
 - configuring for external key
 - management 220
 - for external security key management,
 - configuring 210
 - initial automatic discovery 34
 - introduction 1
 - IP addresses 7
 - manual discovery 36
 - mapping new WWPNs for AIX and
 - Linux 179
 - naming 37
 - setting passwords 36
 - static TCP/IP addresses 9
 - tuning 95, 96, 97, 98, 99, 100
 - tuning options available 95
- Storport Miniport 124
- Subsystem Management window
 - elements 19, 20
 - event log 245
 - Hardware tab 25
 - Host Mappings tab 24
 - multiple software versions 26
 - online help xiv
 - opening 20
 - Setup tab 26
 - Storage and Copy Services tab 21
 - Summary tab 21
- Summary tab
 - See* Subsystem Management window
- support
 - multipath driver 115
 - notifications xviii
 - obtaining xvi, xvii
 - websites xv, xviii
- support notifications xviii
 - receiving xviii
- switch environment 119
- switches
 - in a SAN-attached configuration 6
 - zoning 6
- System Storage Productivity Center xv

T

- target authentication, changing for
 - iSCSI 40
- target discovery, changing for iSCSI 40
- target identification, changing for
 - iSCSI 40

- Task Assistant
 - description 56
 - shortcuts 56
- TCP/IP
 - IPv6 42
- TCP/IP addresses, static
 - assigning to storage subsystems 9
- terminology, FDE 209
- Tivoli Key Lifecycle Manager
 - See* IBM Tivoli Key Lifecycle Manager
- TKLM
 - See* IBM Tivoli Key Lifecycle Manager
- tools
 - lsslot 176
 - PCI Hotplug 175
- trademarks 319
- transfer rate 95
- troubleshooting 245
 - critical events 245
 - Diagnostic Data Capture 263

U

- uninstallation
 - DS TKLM Proxy Code on Windows,
 - for external security key
 - management 215
 - Storage Manager 33
- universal transport mechanism
 - See* UTM device
- updates
 - receiving xviii
- updates (product updates) xviii
- utilities
 - hot_add 162
 - SMdevices 159
 - SMrepassist 163
- UTM device 161

V

- VDS/VSS provider 309
- Veritas 129
 - DMP 156
 - Dynamic Multipathing (DMP) 148
 - File System 156
 - Storage Foundation 167
 - Storage Foundation 5.0 167
 - VolumeManager 148, 156
- Veritas DMP drivers 143
- Veritas DMP DSM 129
- Veritas Storage Foundation
 - LVM scan, disabling for SUSE Linux
 - Enterprise Server 167
 - RDAC module, enabling on RHEL for
 - Storage Foundation 5.0 167
- Veritas Storage Foundation 5.0
 - RDAC module, enabling 167
 - RDAC module, unloading 167
- VMware ESX Server 285
 - cross connections 289
 - mapping LUNs to a partition 289
- VolumeCopy 88

W

- websites
 - documentation xiii
 - FDE best practices 242
 - list xv
 - notification xviii
 - services xvii
 - Solaris failover driver 156
 - support xvii, xviii
 - VMware 289
- window
 - Script Editor 100
- worldwide port name
 - See* WWPN
- write-caching
 - enabling 98
- WWPN
 - mapping to storage subsystem on AIX
 - and Linux 179

Z

- zoning 119
- zoning switches 6



Part Number: 46W6899

Printed in USA

GA32-2221-03



(1P) P/N: 46W6899

