

IBM TotalStorage DS8000



User's Guide

IBM TotalStorage DS8000



User's Guide

Note:

Before using this information and the product it supports, read the information under **Safety and environmental notices** and **Notices**.

Fourth Edition (April 2005)

© Copyright International Business Machines Corporation 2004, 2005. All rights reserved.

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 guide	xi
Safety and environmental notices	xiii
Safety notices	xiii
Environmental notices	xiii
Product recycling	xiii
Disposing of products	xiii
Conventions used in this guide	xiii
Related information	xiv
DS8000 series library	xiv
Other IBM publications	xv
Ordering IBM publications	xx
Web sites	xx
How to send your comments.	xxi
 Summary of Changes for SC26-7623-03 IBM TotalStorage DS8000 User's Guide	 xxiii
 Chapter 1. Introduction to IBM TotalStorage DS8000 series	 1
DS8000 models	1
DS8100 (Model 921)	4
DS8300 (Models 922 and 9A2)	5
Model comparison	5
DS8000 physical footprint	6
DS8000 performance features.	8
Interfaces of the DS8000	8
IBM TotalStorage DS Storage Manager	9
DS command-line interface	9
DS Open application programming interface	9
DS8000 hardware specifics	10
Storage unit structure	10
IBM TotalStorage Management Console.	11
Host systems that the storage unit supports	11
Processor memory	15
Data management features	15
RAID	15
Arrays across loops	15
Storage System LPARs (logical partitions)	15
Copy Services	16
FlashCopy	18
Disaster recovery using Copy Services	19
Comparison of licensed functions	20
Parallel Access Volumes	21
DS8000 limitations	21
Planning data migration.	22
How to select a data migration method	22
 Chapter 2. Installing the DS8000 Storage Manager interface	 25

	Installing the IBM TotalStorage DS Storage Manager on the Windows operating system using the graphical mode	26
	Installing the IBM TotalStorage DS Storage Manager on the Windows operating system in unattended (silent) mode	37
	Removing the IBM TotalStorage DS Storage Manager from a Windows operating system	42
	Removing the DS8000 Storage Manager from a Windows operating system using the silent mode.	44
I	Upgrading the DS Storage Manager on a Windows operating system using the graphical mode	44
I	Upgrading the IBM TotalStorage DS Storage Manager on the Windows operating system in unattended (silent) mode	51
I	Chapter 3. Installing the IBM TotalStorage DS8000 Storage Manager on the Linux operating system	53
	Installing the IBM TotalStorage DS Storage Manager on a Linux operating system using the graphical mode	53
	Installing the IBM TotalStorage DS8000 Storage Manager on the Linux operating system in unattended (silent) mode.	63
	Removing the IBM TotalStorage DS Storage Manager from a Linux operating system	70
	Removing the DS8000 Storage Manager from a Linux operating system using the silent mode	71
I	Upgrading the IBM TotalStorage DS Storage Manager on a Linux operating system	72
I	Upgrading the IBM TotalStorage DS8000 Storage Manager on the Linux operating system in unattended (silent) mode.	79
I	Chapter 4. DS Storage Manager postinstallation instructions	81
	Internet browser support	81
	Stop and start the DS Storage Manager Servers	81
	Stopping and restarting the DS Storage Manager servers on a Windows operating system	81
	Stopping and restarting the DS Storage Manager servers on a Linux operating system	82
	Starting the DS Storage Manager program.	82
	Starting the DS8000 Storage Manager on a Windows operating system	82
	Starting the DS8000 Storage Manager on a Linux operating system	82
	Getting started with the DS Storage Manager console	83
	Activating licenses	84
	Chapter 5. IBM TotalStorage Management Console	87
	Logging onto the Management Console.	87
	Logging off the Management Console	87
	Changing the Management Console password	88
	Accessing the DS Storage Manager	88
	Chapter 6. Administering user accounts	91
	Adding user accounts	91
	Modifying user accounts	91
	Removing user accounts	92
	Unlocking a user account	92
	Unlocking an administrative password	93
	Chapter 7. Storage system logical partition	95
	DS8300 LPAR implementation	95

Required hardware components to build a storage image	96
I/O resources	97
Processor and memory allocation	97
RIO-G interconnect separation	98
Chapter 8. Optional feature codes for the DS8000	99
Using parallel access volumes for S/390 and zSeries hosts	99
Parallel access volumes	100
PAV requirements for the DS8000	100
Overview of the input/output configuration program for the DS8000	101
Guidelines for defining the logical control unit for the DS8000	103
Reconfiguring an LCU with larger capacity 3390 volumes	104
Increasing the number of 3390s formatted within an LSS	104
Copy Services.	104
Copy Services overview	104
IBM TotalStorage Resiliency Family	105
Disaster recovery using Copy Services	106
Copy Services functions across a 2105 and 2107	107
FlashCopy	110
Reverse FlashCopy relationships	111
FlashCopy to Metro Mirror source volumes	112
Multiple relationship FlashCopy	113
Persistent FlashCopy	113
Re-synchronizing FlashCopy volume pairs	114
Creating a Remote FlashCopy operation (Inband FlashCopy on the ESS 2105)	114
Path connections.	116
Remote Mirror and Copy	116
Global Copy	117
Global Mirror	118
Metro Mirror	121
Consistency groups.	122
Failover and failback operations	123
z/OS Copy Services	124
z/OS Global Mirror	124
z/OS Metro/Global Mirror.	125
Chapter 9. Migrating data to a storage unit	127
How to select a data migration method	127
Considerations for replacing existing storage	128
Reformatting open systems host DDMs	128
Overview of the logical volume management software	129
Copying a complete logical volume	129
Using mirroring to migrate data	130
Using the migratepv command	131
Using the direct copy method	132
Using backup and restore methods to migrate data	133
Migrating data from an eServer i5 400 or iSeries host system to the DS8000	134
Migrating data from an S/390 or zSeries host system to the DS8000 storage unit	134
Data migration from a UNIX-based host system to a DS8000 storage unit	135
Migrating data to a DS8000 storage unit using volume management methods	136
Migrating data using the direct copy method	137
Migrating data using backup and restore methods	137

Migrating data from a Windows 2000 or Windows Server 2003-based host system	137
Migrating data between storage units	138
Moving workloads when migrating data between storage units	138
Offloading workloads when migrating data between storage units	138
Chapter 10. Disaster recovery using Copy Services.	141
Performing failover and subsequent failback operations	141
Moving production to Site B after planned outages (failover).	141
Returning production to local site after planned outages	143
Moving production to Site B after unplanned outages (failover).	145
Using Global Mirror for disaster recovery	146
Setting up your environment to use Global Mirror.	147
Using Global Mirror for a planned failover and failback	148
Using Global Mirror for an unplanned failover and failback	149
Recovering with and without the fast reverse restore feature.	151
Appendix. Troubleshooting	153
Analyzing generic alert traps	153
Analyzing service information messages (SIMs) for S/390 and zSeries systems	153
Understanding problem status designations	154
Managing product-specific alert traps	154
Obtaining services for a storage unit	155
Analyzing normal operation problems	155
Managing informational messages from the storage unit	156
Managing the storage unit error messages	156
Accessibility	159
Notices	161
Terms and conditions for downloading and printing publications	162
Trademarks.	163
Electronic emission notices	164
Federal Communications Commission (FCC) statement	164
Industry Canada compliance statement	164
European community compliance statement.	164
Japanese Voluntary Control Council for Interference (VCCI) class A statement	165
Korean Ministry of Information and Communication (MIC) statement.	165
Taiwan class A compliance statement	166
Glossary	167
.	167
Index	195

Figures

1.	A base model (with front cover off) and its main components	2
2.	An expansion model (with front cover off) and its main components	3
3.	Maximum configuration for Model 921	4
4.	Maximum configuration for Models 922 and 9A2.	5
5.	DS8000 physical footprint. Dimensions are in centimeters (inches)..	7
6.	DS8300 LPAR resource allocation	96
7.	Storage image resource allocation in the processor complexes of the DS8300	97
8.	Refreshing target volume — reverse restore	112

Tables

1.	DS8000 model comparison chart	6
2.	Comparison of licensed functions	20
I 3.	Failover and failback implementation	142
I 4.	Failover and failback implementation	145

About this guide

IBM TotalStorage DS8000 is a high-performance, high-capacity series of disk storage that is designed to support continuous operations. This publication provides information that allows you to effectively use the DS8000 storage unit.

Included among this information are procedures and guidelines for performing the following tasks:

- Installing the IBM TotalStorage DS8000 Storage Manager on a Windows or Linux operating system
- Using the Management Console to access the DS Storage Manager
- Using the optional features supported on the storage unit
- Migrating data to or between storage units
- Handling disaster recovery
- Analyzing problems

Customers or service representatives who install, set up, and use the DS8000 storage unit should use this publication.

Safety and environmental notices

This section contains information about safety notices that are used in this guide and environmental notices for this product.

Safety notices

Use this process to find information about safety notices.

To find the translated text for a danger or caution notice:

1. Look for the identification number at the end of each danger notice or each caution notice. In the following examples, the numbers **1000** and **1001** are the identification numbers.

DANGER

A danger notice indicates the presence of a hazard that has the potential of causing death or serious personal injury.

1000

CAUTION:

A caution notice indicates the presence of a hazard that has the potential of causing moderate or minor personal injury.

1001

2. Find the number that matches in the *IBM TotalStorage Solutions Safety Notices for IBM Versatile Storage Server and IBM TotalStorage Enterprise Storage Server*, GC26-7229.

Environmental notices

This section identifies the environmental guidelines that pertain to this product.

Product recycling

This unit contains recyclable materials.

Recycle these materials at your local recycling sites. Recycle the materials according to local regulations. In some areas, IBM provides a product take-back program that ensures proper handling of the product. Contact your IBM representative for more information.

Disposing of products

This topic contains information about how to dispose of products.

This unit might contain batteries. Remove and discard these batteries, or recycle them, according to local regulations.

Conventions used in this guide

The following typefaces are used to show emphasis:

boldface

Text in **boldface** represents menu items and lowercase or mixed-case command names.

italics Text in *italics* is used to emphasize a word. In command syntax, it is used for variables for which you supply actual values.

monospace

Text in monospace identifies the data or commands that you type, samples of command output, or examples of program code or messages from the system.

Related information

The tables in this section list and describe the following publications:

- The publications that make up the IBM® TotalStorage™ DS8000 series library
- Other IBM publications that relate to the DS8000 series
- Non-IBM publications that relate to the DS8000 series

See “Ordering IBM publications” on page xx for information about how to order publications in the IBM TotalStorage DS8000 series publication library. See “How to send your comments” on page xxi for information about how to send comments about the publications.

DS8000 series library

These customer publications make up the DS8000 series library.

Unless otherwise noted, these publications are available in Adobe portable document format (PDF) on a compact disc (CD) that comes with the storage unit. If you need additional copies of this CD, the order number is SK2T-8803. These publications are also available as PDF files by clicking on the **Documentation link** on the following Web site:

<http://www-1.ibm.com/servers/storage/support/disk/ds8100/index.html>

See “Ordering IBM publications” on page xx for information about ordering these and other IBM publications.

Title	Description	Order Number
<i>IBM® TotalStorage® DS: Command-Line Interface User's Guide</i>	This guide describes the commands that you can use from the command-line interface (CLI) for managing your DS8000 configuration and Copy Services relationships. The CLI application provides a set of commands that you can use to write customized scripts for a host system. The scripts initiate predefined tasks in a Copy Services server application. You can use the CLI commands to indirectly control Remote Mirror and Copy and FlashCopy® configuration tasks within a Copy Services server group.	SC26-7625 (See Note.)
<i>IBM TotalStorage DS8000: Host Systems Attachment Guide</i>	This guide provides guidelines for attaching the DS8000 to your host system and for migrating to fibre-channel attachment from a small computer system interface.	SC26-7628 (See Note.)
<i>IBM TotalStorage DS8000: Introduction and Planning Guide</i>	This guide introduces the DS8000 product and lists the features you can order. It also provides guidelines for planning the installation and configuration of the storage unit.	GC35-0495

Title	Description	Order Number
<i>IBM TotalStorage Multipath Subsystem Device Driver User's Guide</i>	This publication describes how to use the IBM Subsystem Device Driver (SDD) on open-systems hosts to enhance performance and availability on the DS8000. SDD creates redundant paths for shared logical unit numbers. SDD permits applications to run without interruption when path errors occur. It balances the workload across paths, and it transparently integrates with applications.	SC30-4096
<i>IBM TotalStorage DS8000: User's Guide</i>	This guide provides instructions for setting up and operating the DS8000 and for analyzing problems.	SC26-7623 (See Note.)
<i>IBM TotalStorage DS Application Programming Interface Reference</i>	This publication provides reference information for the IBM TotalStorage DS application programming interface (API) and provides instructions for installing the Common Information Model Agent, which implements the API.	GC35-0493
<i>IBM TotalStorage DS8000 Messages Reference</i>	This publication provides explanations of error, information, and warning messages that are issued from the DS8000 user interfaces.	GC26-7659
Note: No hardcopy book is produced for this publication. However, a PDF file is available from the following Web site: http://www-1.ibm.com/servers/storage/support/disk/ds8100/index.html		

Other IBM publications

Other IBM publications contain additional information that is related to the DS product library.

The following list is divided into categories to help you find publications that are related to specific topics. Some of the publications are listed under more than one category. See "Ordering IBM publications" on page xx for information about ordering these and other IBM publications.

Title	Description	Order Number
Data-copy services		
<i>z/OS DFSMS Advanced Copy Services</i>	This publication helps you understand and use IBM Advanced Copy Services functions. It describes three dynamic copy functions and several point-in-time copy functions. These functions provide backup and recovery of data if a disaster occurs to your data center. The dynamic copy functions are peer-to-peer remote copy, extended remote copy, and coupled extended remote copy. Collectively, these functions are known as remote copy. FlashCopy, SnapShot, and concurrent copy are the point-in-time copy functions.	SC35-0428
<i>IBM Enterprise Storage Server</i>	This publication, from the IBM International Technical Support Organization, introduces the Enterprise Storage Server and provides an understanding of its benefits. It also describes in detail the architecture, hardware, and functions, including the advanced copy functions, of the Enterprise Storage Server.	SG24-5465
<i>Implementing Copy Services On S/390</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on an Enterprise Storage Server that is attached to an S/390 or zSeries host system. Copy Services functions include peer-to-peer remote copy, extended remote copy, FlashCopy®, and concurrent copy. This publication describes the functions, prerequisites, and corequisites and describes how to implement each function into your environment.	SG24-5680

Title	Description	Order Number
<i>IBM TotalStorage ESS Implementing Copy Services in an Open Environment</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on UNIX, Windows NT®, Windows 2000, Sun Solaris, HP-UX, Tru64, OpenVMS, and iSeries host systems. The Copy Services functions that are described include peer-to-peer remote copy and FlashCopy. This publication describes the functions and shows you how to implement them into your environment. It also shows you how to implement these functions in a high-availability cluster multiprocessing environment.	SG24-5757
Fibre channel		
<i>Fibre Channel Connection (FICON) I/O Interface: Physical Layer</i>	This publication provides information about the fibre-channel I/O interface. This book is also available as a PDF file from the following Web site: http://www.ibm.com/servers/resourcelink/	SA24-7172
<i>Fibre Transport Services (FTS): Physical and Configuration Planning Guide</i>	This publication provides information about fibre-optic and ESCON-trunking systems.	GA22-7234
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide</i>	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08.	SC26-7350
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide</i>	This guide describes the IBM SAN Fibre Channel Switch and the IBM TotalStorage ESS Specialist. It provides information about the commands and how to manage the switch with Telnet and the Simple Network Management Protocol.	SC26-7349
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide</i>	This publication describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S16. It is intended for trained service representatives and service providers.	SC26-7352
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide</i>	This guide introduces the IBM SAN Fibre Channel Switch 2109 Model S16 and tells you how to manage and monitor the switch using zoning and how to manage the switch remotely.	SC26-7351
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the Enterprise Storage Server. It provides you with a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment.	SG24-6113
Open-systems hosts		
<i>ESS Solutions for Open Systems Storage: Compaq AlphaServer, HP, and Sun</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure the Enterprise Storage Server when you attach Compaq AlphaServer (running Tru64 UNIX), HP, and Sun hosts. This book does not cover Compaq AlphaServer that is running the OpenVMS operating system. This book also focuses on the settings that are required to give optimal performance and on the settings for device driver levels. This book is for the experienced UNIX professional who has a broad understanding of storage concepts.	SG24-6119

Title	Description	Order Number
<i>IBM TotalStorage ESS Implementing Copy Services in an Open Environment</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on UNIX or Windows 2000 host systems. The Copy Services functions that are described include peer-to-peer remote copy and FlashCopy. This publication describes the functions and shows you how to implement them into your environment. It also shows you how to implement these functions in a high-availability cluster multiprocessing environment.	SG24-5757
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the Enterprise Storage Server. It gives you a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment.	SG24-6113
S/390 and zSeries hosts		
<i>Device Support Facilities: User's Guide and Reference</i>	This publication describes the IBM Device Support Facilities (ICKDSF) product that are used with IBM direct access storage device (DASD) subsystems. ICKDSF is a program that you can use to perform functions that are needed for the installation, the use, and the maintenance of IBM DASD. You can also use it to perform service functions, error detection, and media maintenance.	GC35-0033
<i>z/OS Advanced Copy Services</i>	This publication helps you understand and use IBM Advanced Copy Services functions. It describes three dynamic copy functions and several point-in-time copy functions. These functions provide backup and recovery of data if a disaster occurs to your data center. The dynamic copy functions are peer-to-peer remote copy, extended remote copy, and coupled extended remote copy. Collectively, these functions are known as remote copy. FlashCopy, SnapShot, and concurrent copy are the point-in-time copy functions.	SC35-0428
<i>DFSMS/MVS V1: Remote Copy Guide and Reference</i>	This publication provides guidelines for using remote copy functions with S/390 and zSeries hosts.	SC35-0169
<i>Fibre Transport Services (FTS): Physical and Configuration Planning Guide</i>	This publication provides information about fibre-optic and ESCON-trunking systems.	GA22-7234
<i>Implementing ESS Copy Services on S/390</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on an Enterprise Storage Server that is attached to an S/390 or zSeries host system. Copy Services functions include peer-to-peer remote copy, extended remote copy, FlashCopy, and concurrent copy. This publication describes the functions, prerequisites, and corequisites and describes how to implement each function into your environment.	SG24-5680
<i>ES/9000, ES/3090: IOCP User Guide Volume A04</i>	This publication describes the Input/Output Configuration Program that supports the Enterprise Systems Connection (ESCON) architecture. It describes how to define, install, and configure the channels or channel paths, control units, and I/O devices on the ES/9000 processors and the IBM ES/3090 Processor Complex.	GC38-0097
<i>IOCP User's Guide, IBM e(logo)server zSeries 800 and 900</i>	This publication describes the Input/Output Configuration Program that supports the zSeries 800 and 900 servers. This publication is available in PDF format by accessing ResourceLink at the following Web site: www.ibm.com/servers/resourceLink/	SB10-7029

Title	Description	Order Number
<i>IOCP User's Guide, IBM e(logo)server zSeries</i>	This publication describes the Input/Output Configuration Program that supports the zSeries server. This publication is available in PDF format by accessing ResourceLink at the following Web site: www.ibm.com/servers/resourceLink/	SB10-7037
<i>S/390: Input/Output Configuration Program User's Guide and ESCON Channel-to-Channel Reference</i>	This publication describes the Input/Output Configuration Program that supports ESCON architecture and the ESCON multiple image facility.	GC38-0401
<i>IBM z/OS Hardware Configuration Definition User's Guide</i>	This guide provides conceptual and procedural information to help you use the z/OS Hardware Configuration Definition (HCD) application. It also explains: <ul style="list-style-type: none"> • How to migrate existing IOCP/MVSCP definitions • How to use HCD to dynamically activate a new configuration • How to resolve problems in conjunction with MVS/ESA HCD 	SC33-7988
<i>OS/390: Hardware Configuration Definition User's Guide</i>	This guide provides detailed information about the input/output definition file and about how to configure parallel access volumes. This guide discusses how to use Hardware Configuration Definition for both OS/390® and z/OS V1R1.	SC28-1848
<i>OS/390 V2R10.0: MVS System Messages Volume 1 (ABA - ASA)</i>	This publication lists OS/390 MVS™ system messages ABA to ASA.	GC28-1784
<i>Using IBM 3390 Direct Access Storage in a VM Environment</i>	This publication provides device-specific information for the various models of the 3390 and describes methods you can use to manage storage efficiently using the VM operating system. It provides guidance on managing system performance, availability, and space through effective use of the direct access storage subsystem.	GG26-4575
<i>Using IBM 3390 Direct Access Storage in a VSE Environment</i>	This publication helps you use the 3390 in a VSE environment. It includes planning information for adding new 3390 units and instructions for installing devices, migrating data, and performing ongoing storage management activities.	GC26-4576
<i>Using IBM 3390 Direct Access Storage in an MVS Environment</i>	This publication helps you use the 3390 in an MVS environment. It includes device-specific information for the various models of the 3390 and illustrates techniques for more efficient storage management. It also offers guidance on managing system performance, availability, and space utilization through effective use of the direct access storage subsystem.	GC26-4574
<i>z/Architecture Principles of Operation</i>	This publication provides a detailed definition of the z/Architecture™. It is written as a reference for use primarily by assembler language programmers and describes each function at the level of detail needed to prepare an assembler language program that relies on a particular function. However, anyone concerned with the functional details of z/Architecture will find this publication useful.	SA22-7832
SAN		

Title	Description	Order Number
<i>IBM OS/390 Hardware Configuration Definition User's Guide</i>	<p>This guide explains how to use the Hardware Configuration Data application to perform the following tasks:</p> <ul style="list-style-type: none"> • Define new hardware configurations • View and modify existing hardware configurations • Activate configurations • Query supported hardware • Maintain input/output definition files (IODFs) • Compare two IODFs or compare an IODF with an actual configuration • Print reports of configurations • Create graphical reports of a configuration • Migrate existing configuration data 	SC28-1848
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide</i>	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08.	SC26-7350
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide</i>	This guide describes the IBM SAN Fibre Channel Switch and the IBM TotalStorage ESS Specialist. It provides information about the commands and how to manage the switch with Telnet and the Simple Network Management Protocol (SNMP).	SC26-7349
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide</i>	This publication describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S16. It is intended for trained service representatives and service providers.	SC26-7352
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide</i>	This guide introduces the IBM SAN Fibre Channel Switch 2109 Model S16 and tells you how to manage and monitor the switch using zoning and how to manage the switch remotely.	SC26-7351
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the Enterprise Storage Server. It provides you with a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment.	SG24-6113
Seascope family		
<i>IBM Enterprise Storage Server</i>	This publication, from the IBM International Technical Support Organization, introduces the Enterprise Storage Server and provides an understanding of its benefits. It also describes in detail the architecture, hardware, and functions, including the advanced copy functions, of the Enterprise Storage Server.	SG24-5465
<i>IBM Enterprise Storage Server Performance Monitoring and Tuning Guide</i>	This guide, from the IBM International Technical Support Organization, provides guidance on the best way to configure, monitor, and manage your Enterprise Storage Server to ensure optimum performance.	SG24-5656
<i>IBM Versatile Storage Server: Introduction and Planning Guide</i>	This publication introduces the IBM Versatile Storage Server™ and lists the features you can order. It also provides planning information for both 2105 Models B09 and 100.	GC26-7223
<i>Implementing the IBM Enterprise Storage Server in Your Environment</i>	This publication, from the IBM International Technical Support Organization, can help you install, tailor, and configure the Enterprise Storage Server in your environment.	SG24-5420
Storage management		

Title	Description	Order Number
<i>Device Support Facilities: User's Guide and Reference</i>	This publication describes the IBM Device Support Facilities (ICKDSF) product used with IBM direct access storage device (DASD) subsystems. ICKDSF is a program that you can use to perform functions that are needed for the installation, the use, and the maintenance of IBM DASD. You can also use it to perform service functions, error detection, and media maintenance.	GC35-0033
<i>IBM TotalStorage Solutions Handbook</i>	This handbook, from the IBM International Technical Support Organization, helps you understand what makes up enterprise storage management. The concepts include the key technologies that you must know and the IBM subsystems, software, and solutions that are available today. It also provides guidelines for implementing various enterprise storage administration tasks so that you can establish your own enterprise storage management environment.	SG24-5250

Ordering IBM publications

This section tells you how to order copies of IBM publications and how to set up a profile to receive notifications about new or changed publications.

IBM publications center

The publications center is a worldwide central repository for IBM product publications and marketing material.

The IBM publications center offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download free of charge. You can also order publications. The publications center displays prices in your local currency. You can access the IBM publications center through the following Web site:

<http://www.ibm.com/shop/publications/order>

Publications notification system

The IBM publications center Web site offers you a notification system for IBM publications.

If you register, you can create your own profile of publications that interest you. The publications notification system sends you a daily e-mail that contains information about new or revised publications that are based on your profile.

If you want to subscribe, you can access the publications notification system from the IBM publications center at the following Web site:

<http://www.ibm.com/shop/publications/order>

Web sites

The following Web sites provide information about the IBM TotalStorage DS8000 series and other IBM storage products.

Type of Storage Information	Web Site
Concurrent Copy for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/

Type of Storage Information	Web Site
Copy Services command-line interface (CLI)	http://www-1.ibm.com/servers/storage/support/software/cscli.html
DS8000 series publications	http://www-1.ibm.com/servers/storage/support/disk/ds8100/index.html Click Documentation .
FlashCopy for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
Host system models, operating systems, and adapters that the storage unit supports	http://www.ibm.com/servers/storage/disk/ds8000/interop.html Click Interoperability matrix .
IBM Disk Storage Feature Activation (DSFA)	http://www.ibm.com/storage/dsfa
IBM storage products	http://www.storage.ibm.com/
IBM TotalStorage DS8000 series	http://www-1.ibm.com/servers/storage/disk/ds8000
IBM version of the Java (JRE) that is often required for IBM products	http://www-106.ibm.com/developerworks/java/jdk/
Multiple Device Manager (MDM)	http://www.ibm.com/servers/storage/support/ Click Storage Virtualization .
Remote Mirror and Copy (formerly PPRC) for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
SAN fibre channel switches	http://www.ibm.com/storage/fcswitch/
Storage Area Network Gateway and Router	http://www-1.ibm.com/servers/storage/support/san/index.html
Subsystem Device Driver (SDD)	http://www-1.ibm.com/servers/storage/support/software/sdd.html
z/OS Global Mirror (formerly XRC) for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this information or any other DS8000 series documentation, you can submit them in the following ways:

- e-mail

Submit your comments electronically to the following e-mail address:

starpubs@us.ibm.com

Be sure to include the name and order number of the book and, if applicable, the specific location of the text you are commenting on, such as a page number or table number.

- Mail

Fill out the Readers' Comments form (RCF) at the back of this book. Return it by mail or give it to an IBM representative. If the RCF has been removed, you can address your comments to:

International Business Machines Corporation
RCF Processing Department

Department 61C
9032 South Rita Road
TUCSON AZ 85775-4401

Summary of Changes for SC26-7623-03 IBM TotalStorage DS8000 User's Guide

This document contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change. This summary of changes describes new functions that have been added to this release.

New Information

This edition includes the following new information:

- Chapter 2, the section *Upgrading the DS Storage Manager on Windows, in silent mode* was added to show the steps for upgrading the DS Storage Manager.
- Chapter 3, the section *Upgrading the DS Storage Manager on Linux, in silent mode* was added to show the steps for upgrading the DS Storage Manager.
- Chapter 8, the section *FlashCopy* information was added to explain how to set up a Remote FlashCopy operation, which allows a source volume of a remote mirror and copy volume pair on a source (local) storage unit to be sent across paths (acting as a conduit) to a target storage unit to enable a FlashCopy pair to be established at the target site. This eliminates the need for a network connection to the target site solely for the management of FlashCopy.

Changed Information

The following information was updated or changed:

- Chapter 2, the section *Installing the DS Storage Manager on the Windows operating system using the graphical mode* was updated to include information about new version checking and progress windows.
- Chapter 3, the section *Installing the DS Storage Manager on the Linux operating system using the graphical mode* was updated to include information about new version checking and progress windows.

Chapter 1. Introduction to IBM TotalStorage DS8000 series

IBM TotalStorage DS8000 is a high-performance, high-capacity series of disk storage that is designed to support continuous operations. DS8000 series models (machine type 2107) use the IBM POWER5™ server technology that is integrated with the IBM Virtualization Engine™ technology. DS8000 series models consist of a storage unit and one or two management consoles, two being the recommended configuration. The graphical user interface (GUI) or the command-line interface (CLI) allows you to logically partition storage and use the built-in Copy Services functions. For high-availability, hardware components are redundant.

You can read the following information to familiarize yourself with the DS8000 series:

- DS8000 series models, including a model comparison chart
- DS8000 series performance features
- DS8000 series interfaces
- DS8000 series hardware specifics
- Supported systems for open systems, S/390®, and zSeries® hosts
- Data management elements
- Copy Services
- Data migration

DS8000 models

The DS8000 series offers various choices of base and expansion models, so you can configure storage units that meet your performance and configuration needs.

DS8100

The DS8100 (Model 921) features a dual two-way processor complex and support for one expansion frame.

DS8300

The DS8300 (Models 922 and 9A2) features a dual four-way processor complex and support for one or two expansion frames. The Model 9A2 supports two IBM TotalStorage System LPARs (Logical Partitions) in one storage unit.

The DS8000 expansion frames (Models 92E and 9AE) expand the capabilities of the base models. You can attach the Model 92E to either the Model 921 or the Model 922 to expand their capabilities. You can attach the Model 9AE to expand the Model 9A2.

Figure 1 on page 2 provides a high-level view at the components for the DS8100 and DS8300 base models (Models 921, 922, and 9A2).





Figure 1. A base model (with front cover off) and its main components

The following notes provide additional information about the labeled components on Figure 1:

1. The rack power area of the base models provides redundant power supplies (two primary power supplies), power control cards, and backup battery assemblies to help protect data in the event of a loss of external power. Model 921 contains two batteries. Model 922 and Model 9A2 contain three batteries to support the 4-way processors.
2. All base models allow up to eight disk enclosures, which contain the disk drives. In a maximum configuration, each base model can hold up to 128 disk drives.
3. All base models contain one Management Console (MC), a keyboard, and display. The Management Console is the focal point for configuration, copy services management, and maintenance activities.
4. All base models contain two processor enclosures. The processor enclosures contain the computer electronic complex (CEC). The Model 921 processor enclosures have 2-way processors. Processor enclosures on the Model 922 and Model 9A2 have 4-way processors.
5. All base models contain I/O enclosures and adapters. The I/O enclosures hold the adapters and provide connectivity between the adapters and the processors. Both device adapters and host adapters are installed in the I/O enclosure.

Figure 2 on page 3 provides a high-level view at the components for the expansion models (Models 92E and 9AE).



1vxc86

Figure 2. An expansion model (with front cover off) and its main components

The following notes provide additional information about the labeled components on Figure 2:

1. The rack power area of each expansion model provides redundant power supplies (two primary power supplies) and power control cards. If the expansion unit contains I/O enclosures or the extended power line disturbance (PLD) feature, three backup battery assemblies are also installed to help protect data in the event of a loss of external power.
2. All expansion models allow up to 16 disk enclosures, which contain the disk drives. In a maximum configuration, each expansion model can hold 256 disk drives.
3. Expansion models can contain I/O enclosures and adaptors if they are the first expansion units that are attached to either a Model 922 or a Model 9A2. The second expansion model in a 922 or 9A2 configuration cannot have I/O enclosures and adaptors, nor can any expansion unit that is attached to a Model 921. If the expansion unit contains I/O enclosures, the enclosures provide connectivity between the adaptors and the processors. The adaptors contained in the I/O enclosures can be either device or host adaptors.

DS8100 (Model 921)

The IBM TotalStorage DS8100, which is Model 921, offers many features.

These features include:

- Dual two-way processor complex
- Up to 128 disk drives, for a maximum capacity of 38.4 TB
- Up to 128 GB of processor memory (cache)
- Up to 16 fibre-channel/FICON or ESCON® host adapters

The DS8100 model can support one expansion frame. With one expansion frame, you can expand the capacity of the Model 921 as follows:

- Up to 384 disk drives, for a maximum capacity of 115.2 TB

Note: IBM service representatives can upgrade a Model 921 in the field when you order a model conversion to a Model 922 or Model 9A2.

Figure 3 shows the maximum configuration of a Model 921, which is the 921 base model plus one 92E expansion model.




Figure 3. Maximum configuration for Model 921

DS8300 (Models 922 and 9A2)

IBM TotalStorage DS8300 models (Model 922 and Model 9A2) offer higher performance and capacity than the DS8100. The Model 9A2 also enables you to create two storage system LPARs (or images) within the same storage unit.

Both DS8300 models offer the following features:

- Dual four-way processor complex
- Up to 128 disk drives, for a maximum capacity of 38.4 TB
- Up to 256 GB of processor memory (cache)
- Up to 16 fibre-channel/FICON or ESCON host adapters

The DS8300 models can support either one or two expansion frames. With expansion frames, you can expand the Model 922 and 9A2 as follows:

- With one expansion frame, you can support the following expanded capacity and number of adapters:
 - Up to 384 disk drives, for a maximum capacity of 115.2 TB
 - Up to 32 fibre-channel/FICON or ESCON host adapters
- With two expansion frames, you can support the following expanded capacity:
 - Up to 640 disk drives, for a maximum capacity of 192 TB

Figure 4 shows the maximum configuration of a Model 922 and 9A2. Both of these models can attach up to two expansion models. Model 922 uses the 92E expansion model. Model 9A2 uses the 9AE expansion model.

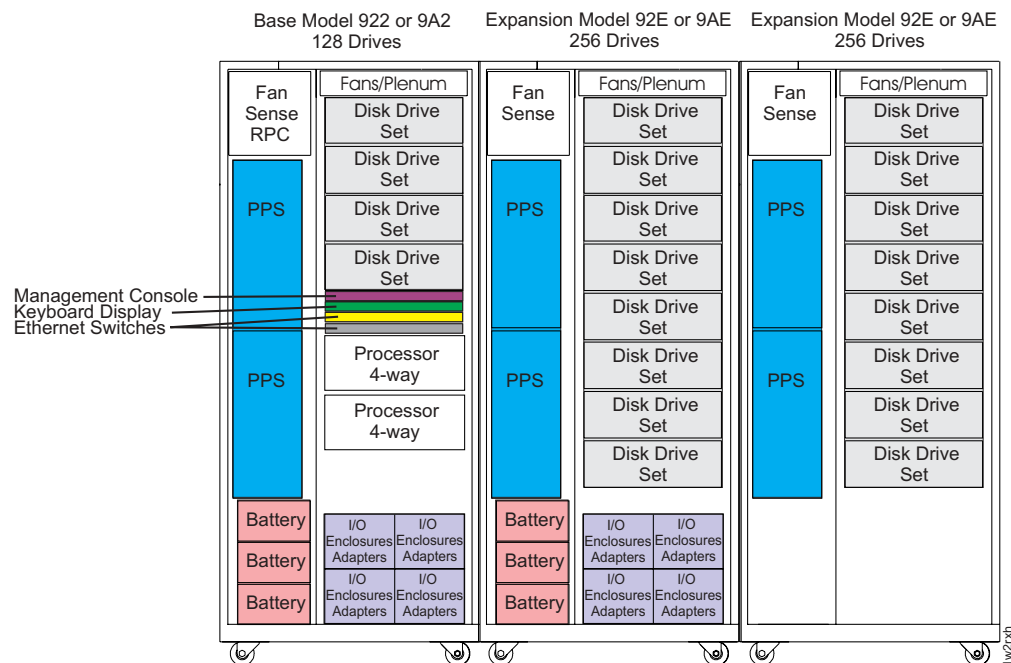


Figure 4. Maximum configuration for Models 922 and 9A2

Model comparison

DS8000 models vary on the processor types that they use, their disk capacity, and other supported features.

Table 1 compares the various supported features that are available on different DS8000 model configurations.

Table 1. DS8000 model comparison chart

Base model	Images (logical partitions)	Expansion models	Processor type	Disk drives	Processor memory	Host adapters
Model 921	1 image	None	2-way	Up to 128 disks (maximum capacity of 38.4 TB)	Up to 128 GB	Up to 16 host adapters
		With 1 expansion model (92E)	2-way	Up to 384 disks (maximum capacity of 115.2 TB)	Up to 128 GB	Up to 16 host adapters
Model 922	1 image	None	4-way	Up to 128 disks (maximum capacity of 38.4 TB)	Up to 256 GB	Up to 16 host adapters
		With 1 expansion model (92E)	4-way	Up to 384 disks (maximum capacity of 115.2 TB)	Up to 256 GB	Up to 32 host adapters
		With 2 expansion models (92E)	4-way	Up to 640 disks (maximum capacity of 192 TB)	Up to 256 GB	Up to 32 host adapters
Model 9A2	2 images	None	4-way	Up to 128 disks (maximum capacity of 38.4 TB)	Up to 256 GB	Up to 16 host adapters
		With 1 expansion model (9AE)	4-way	Up to 384 disks (maximum capacity of 115.2 TB)	Up to 256 GB	Up to 32 host adapters
		With 2 expansion models (9AE)	4-way	Up to 640 disks (maximum capacity of 192 TB)	Up to 256 GB	Up to 32 host adapters

Note: IBM service representatives can upgrade a Model 921 in the field when you order a model conversion to a Model 922 or a Model 9A2.

DS8000 physical footprint

The physical footprint dimensions, caster locations, and cable openings for a DS8000 unit help you plan your installation site.

Figure 5 shows the overall physical footprint of a DS8000 unit.




Figure 5. DS8000 physical footprint. Dimensions are in centimeters (inches).

The following dimensions are labeled on Figure 5:

1. Front cover width
2. Front service clearance
3. Back cover widths
4. Back service clearance
5. Clearance to allow front cover to open
6. Distance between casters
7. Depth of frame without covers
8. Depth of frame with covers
9. Minimum dimension between casters and outside edges of frames

DS8000 performance features

The DS8000 series is designed to provide you with high-performance, connectivity, and reliability allowing your workload to be easily consolidated into a single storage subsystem.

The following list provides an overview of some of the features that are associated with the DS8000 series:

POWER5 processor technology

The DS8000 features IBM POWER5 server technology. Depending on workload, the maximum host I/O operations per second of the DS8100 Model 921 is up to three times the maximum operations per second of the ESS Model 800. The maximum host I/O operations per second of the DS8300 Model 922 or 9A2 is up to six times the maximum of the ESS model 800.

Industry standard fibre channel disk drives

The DS8000 series offers a selection of fibre channel disk drives, including 300 GB drives, allowing a DS8100 to scale up to 115.2 TB of capacity and a DS8300 to scale up to 192 TB of capacity.

Four-port fibre channel/FICON adapters

These new adapters not only enhance connectivity, but increase configuration flexibility because the individual ports can be configured to support fibre channel or FICON.

HOST attachments

The DS8100 offers up to 16 host adapters (64 FCP/FICON ports) and the DS8300 offers up to 32 host adapters (128 FCP/FICON ports), further increasing your ability to share your disk storage.

IBM Standby Capacity on Demand

The IBM Standby Capacity on Demand (Standby CoD) offering allows the installation of inactive disk drives that can be easily activated as business needs require. With this offering, up to four Standby CoD disk drive sets (64 disk drives) can be factory or field installed into your system. To activate, you logically configure the disk drives for use—a non-disruptive activity that does not require intervention from IBM. Upon activation of any portion of the Standby CoD disk drive set, you must place an order with IBM to initiate billing for the activated set. At that time, you can also order replacement Standby CoD disk drive sets.

Online Information Center

The online Information Center is an information database that provides you the opportunity to quickly familiarize yourself with the major aspects of the DS8000 and to easily recognize the topics for which you might require more information. It provides information regarding user assistance for tasks, concepts, reference, user scenarios, tutorials, and other types of user information. Because the information is all in one place rather than across multiple publications, you can access the information that you need more efficiently and effectively.

Interfaces of the DS8000

The DS8000 provides management interfaces that contain management tools and utilities to help you increase productivity.

IBM TotalStorage DS Storage Manager

The IBM TotalStorage DS Storage Manager is a program interface that is used to perform logical configurations and Copy Services management functions.

The DS Storage Manager program is installed as a GUI (graphical mode) or as an unattended (silent mode) installation for the supported operating systems. It can be accessed from any location that has network access using a Web browser. It offers you the following choices:

Simulated configuration

You install this component on your PC or the Master Console which provides the ability to create or modify logical configurations when your storage unit is disconnected from the network. After creating the configuration, you can save it and then apply it to a network attached storage unit at a later time.

Real-time configuration

This component is preinstalled on your HMC. It provides you the ability to create logical configurations and use Copy Services features when your storage unit is attached to the network. This component provides you with real-time (online) configuration support.

DS command-line interface

The IBM TotalStorage DS Command-Line Interface (CLI) enables open systems hosts to invoke and manage FlashCopy and Metro and Global Mirror functions through batch processes and scripts.

The command-line interface provides a full-function command set that allows you to check your storage unit configuration and perform specific application functions when necessary.

The following list highlights a few of the specific types of functions that you can perform with the DS command-line interface:

- Check and verify your storage unit configuration
- Check the current Copy Services configuration that is used by the storage unit
- Create new logical storage and Copy Services configuration settings
- Modify or delete logical storage and Copy Services configuration settings

Note: You cannot install the DS CLI on a Windows 64-bit operating system.

DS Open application programming interface

The DS Open application programming interface (API) is a nonproprietary storage management client application that supports routine LUN management activities, such as LUN creation, mapping and masking, and the creation or deletion of RAID5 and RAID10 volume spaces. The DS Open API also enables Copy Services functions such as FlashCopy and Remote Mirror and Copy (formally known as peer-to-peer remote copy).

The IBM TotalStorage DS Open API helps integrate DS configuration management support into storage resource management (SRM) applications, which allow customers to benefit from existing SRM applications and infrastructures. The DS Open API also enables the automation of configuration management through customer-written applications. Either way, the DS Open API presents another option

for managing storage units by complementing the use of the IBM TotalStorage DS Storage Manager web-based interface and the DS command-line interface.

You must implement the DS Open API through the IBM TotalStorage Common Information Model (CIM) agent, a middleware application that provides a CIM-compliant interface. The DS Open API uses the CIM technology to manage proprietary devices as open system devices through storage management applications. The DS Open API allows these storage management applications to communicate with a storage unit.

DS8000 hardware specifics

The DS8000 models offer a high degree of availability and performance through the use of redundant components that can be replaced while the system is operating. You can use the DS8000 models with a mix of different operating systems and clustered and nonclustered variants of the same operating systems.

Contributing to the high degree of availability and reliability are the structure of the DS8000 storage unit, the host systems it supports, and its processor memory and processor speeds.

Storage unit structure

The design of the storage unit, which contains the base model and the expansion models, contributes to the high degree of availability that is associated with the DS8000. The primary components that support high availability within the storage unit are the storage server, the processor complex, and the rack power control card.

The storage unit also has redundant fibre switches, with fibre fabric connections to both processor complexes. The entire power system at all levels is fully redundant. There are redundant private LAN networks between each storage unit and the management consoles (even if there is only one management console).

Storage unit

A storage unit consists of a storage server and two or more storage (disk) enclosures that are packaged in one or more racks with associated power supplies, batteries, and cooling.

Storage server

A storage server consists of two processor complexes, two or more I/O enclosures, and a pair of rack power control cards.

Processor complex

A processor complex controls and manages the storage unit to perform the function of the storage server. The two processor complexes form a redundant pair such that if either processor complex fails, the remaining processor complex performs all storage server functions.

Rack power control card

A redundant pair of rack power control (RPC) cards coordinate the power management within the storage unit. The RPC cards are attached to the service processors in each processor complex, the primary power supplies in each rack, and indirectly to the fan/sense cards and storage enclosures in each rack.

Multi-path subsystem device driver

The IBM TotalStorage Multi-path subsystem device driver (SDD) provides load balancing and enhanced data availability capability in configurations with more than one I/O path between the host server and the DS8000

series system. Load balancing can help reduce or eliminate I/O bottlenecks that occur when many I/O operations are directed to common devices that are using the same I/O path. The SDD also helps eliminate a potential single point of failure by automatically rerouting I/O operations when a path failure occurs, thereby supporting enhanced data availability capability.

IBM TotalStorage Management Console

The IBM TotalStorage Management Console is the focal point for configuration, copy services management, and maintenance activities.

The Management Console is a dedicated workstation that is physically located (installed) inside your DS8100 and DS8300 and can automatically monitor the state of your system, notifying you and IBM when service is required. The Management Console can also be connected to your network to enable centralized management of your system using the IBM TotalStorage DS Command-Line Interface or the storage management software that uses the IBM TotalStorage DS Open API.

To provide continuous availability of customer access to the management console functions, having a second management console is recommended. This second management console can be provided in two ways:

- **External** (outside the 2107) - This console is generally installed in the customer-provided rack. It uses the same hardware as the internal management console.
- **Internal** (a second management console that resides within a second 2107) - This console can be cabled together with the primary internal management console of the first 2107.'

Host systems that the storage unit supports

The DS8000 storage unit provides a variety of host attachments so that you can consolidate storage capacity and workloads for open-systems hosts, S/390 hosts, and eServer™ zSeries hosts.

The DS8100 Model 921 supports a maximum of 16 host adapters and 4 device adapter pairs, and the DS8300 Models 922 and 9A2 supports a maximum of 32 host adapters and 8 device adapter pairs.

You can configure the storage unit for any of the following system adapter types and protocols:

- Fibre-channel adapters, for support of fibre-channel protocol (FCP) and fibre connection (FICON®) protocol
- Enterprise Systems Connection Architecture® (ESCON) adapters

For fibre-channel attachments, you can establish zones. The zones must contain a single port that is attached to a system adapter with the desired number of ports that are attached to the storage unit. By establishing zones, you reduce the possibility of interactions between system adapters in switched configurations. You can establish the zones by using either of two zoning methods:

- Port number
- Worldwide port name (WWPN)

You can configure switch ports that are attached to the storage unit in more than one zone. This enables multiple system adapters to share access to the storage

unit fibre-channel ports. Shared access to a storage unit fibre-channel port might come from host platforms that support a combination of bus adapter types and the operating systems.

Attaching a DS8000 series to an open-systems host with fibre-channel adapters

You can attach a DS8000 series to an open-systems host with fibre-channel adapters.

Fibre channel is a 1 Gbps or 2 Gbps, full-duplex, serial communications technology to interconnect I/O devices and host systems that are separated by tens of kilometers.

The IBM TotalStorage DS8000 series supports SAN connections at 1 Gbps to 4 Gbps with 2 Gbps host bus adapters. The DS8000 series negotiates automatically and determines whether it is best to run at 1 Gbps link or 2 Gbps link. The IBM TotalStorage DS8000 series detects and operates at the higher link speed.

Fibre channel transfers information between the sources and the users of the information. This information can include commands, controls, files, graphics, video, and sound. Fibre-channel connections are established between fibre-channel ports that reside in I/O devices, host systems, and the network that interconnects them. The network consists of elements like switches, bridges, and repeaters that are used to interconnect the fibre-channel ports.

Fibre-channel (SCSI-FCP) attached open-systems hosts that the storage unit supports

You can attach the DS8000 storage unit to fibre-channel (SCSI-FCP) attached open-systems hosts.

Each storage unit fibre-channel adapter has four ports. Each port has a unique worldwide port name (WWPN). You can configure the port to operate with the SCSI-FCP upper-layer protocol. Shortwave adapters and longwave adapters are available on the storage unit.

Fibre-channel adapters for SCSI-FCP support provide the following configurations:

- A maximum of 64 host ports for DS8100 Model 921 and a maximum of 128 host ports for DS8300 Models 922 and 9A2
- A maximum of 509 host logins per fibre-channel port
- A maximum of 2000 N-port logins per storage image
- Access to all 63.7KB LUNs per target (one target per host adapter), depending on host type
- Either arbitrated loop, switched fabric, or point-to-point topologies

The storage unit supports the following host systems for shortwave fibre-channel attachment and longwave fibre-channel attachment:

- IBM eServer iSeries™ servers that run an IBM OS/400® operating system
- IBM eServer i5 servers that run an IBM i5/OS™ operating system
- IBM RS/6000®, pSeries®, eServer p5, RS/6000 SP™, and pSeries SP server that run an IBM AIX® operating system
- IBM zSeries servers that run the Linux™ operating system
- IBM zSeries servers that run the z/VM® operating system
- Apple Macintosh servers

- Fujitsu PRIMEPOWER servers that run the Solaris operating system
- Hewlett-Packard servers that run an HP Tru64 UNIX® operating system
- Hewlett-Packard servers that run an HP OpenVMS operating system
- Hewlett-Packard servers that run an HP-UX operating system
- Intel-based servers that run a Microsoft® Windows® 2000 operating system
- Intel-based servers that run a Microsoft Windows Server 2003 operating system for both 32-bit and 64-bit configurations
- Intel-based servers that run a Novell NetWare operating system
- IBM servers that run the IBM SAN Volume Controller operating system
- IBM servers that run the IBM SAN File System operating system
- iSCSI Gateway servers that run a Microsoft Windows 2000, Microsoft Windows 2003, or UNIX operating system
- Linux servers that run a Red Hat Linux, Red Flag Linux, Asianux, and SUSE Linux operating system
- SGI servers that run an IRIX operating system
- Sun servers that run a Solaris operating system

See the Interoperability Matrix at

<http://www.ibm.com/servers/storage/disk/ds8000/interop.html> for details about types, models, adapters, and the operating systems that the storage unit supports.

ESCON-attached S/390 and zSeries hosts that the storage unit supports

You can attach the DS8000 storage unit to the ESCON-attached S/390 and zSeries hosts.

With ESCON adapters, the storage unit provides the following configurations:

- A maximum of 32 host ports for DS8100 Model 921 and a maximum of 64 host ports for DS8300 Models 922 and 9A2
- A maximum of 64 logical paths per port
- Access to 16 control-unit images (4096 CKD devices) over a single ESCON port on the storage unit
- Zero to 64 ESCON channels; two per ESCON host adapter
- Two ESCON links with each link that supports up to 64 logical paths

A DS8100 storage unit supports up to 16 host adapters that provide a maximum of 32 ESCON links per machine. A DS8300 storage unit supports up to 32 host adapters that provide a maximum of 64 ESCON links per machine.

Note: ESCON host channels limit the number of devices per channel to 1024. To fully access 4096 devices on a storage unit, it is necessary to connect a minimum of four ESCON host channels to the storage unit. You can access the devices through a switch to a single storage unit ESCON port. This method exposes four control-unit images (1024 devices) to each host channel.

The FICON bridge card in ESCON director 9032 Model 5 enables a FICON bridge channel to connect to ESCON host adapters in the storage unit. The FICON bridge architecture supports up to 16 384 devices per channel. This enables you to attach other control units or other storage units to the same host channel up to the limit that the host supports.

The storage unit supports the following operating systems for S/390 and zSeries hosts:

- Transaction Processing Facility (TPF)
- Virtual Storage Extended/Enterprise Storage Architecture (VSE/ESA™)
- z/OS®
- z/VM
- Linux

For details about models and the operating system versions and releases that the storage unit supports for these host systems, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

FICON-attached S/390 and zSeries hosts that the storage unit supports

You can attach the DS8000 storage unit to FICON-attached S/390 and zSeries hosts.

Each storage unit fibre-channel adapter has four ports. Each port has a unique world wide port name (WWPN). You can configure the port to operate with the FICON upper-layer protocol. When configured for FICON, the fibre-channel port supports connections to a maximum of 128 FICON hosts. On FICON, the fibre-channel adapter can operate with fabric or point-to-point topologies. With fibre-channel adapters that are configured for FICON, the storage unit provides the following configurations:

- Either fabric or point-to-point topologies
- A maximum of 64 host ports for DS8100 Model 921 and a maximum of 128 host ports for DS8300 Models 922 and 9A2
- A maximum of 509 channel connections per fibre-channel port
- A maximum of 2048 logical paths on each fibre-channel port
- Access to all 32 control-unit images (8000 CKD devices) over each FICON port

Note: FICON host channels limit the number of devices per channel to 16 384. To fully access 65 280 devices on a storage unit, it is necessary to connect a minimum of four FICON host channels to the storage unit. You can access the devices through a switch to a single storage unit FICON port. This method exposes 64 control-unit images (16 384 devices) to each host channel.

The storage unit supports the following operating systems for S/390 and zSeries hosts:

- Transaction Processing Facility (TPF)
- Virtual Storage Extended/Enterprise Storage Architecture (VSE/ESA)
- z/OS
- z/VM
- Linux

For details about models, versions of operating systems, and releases that the storage unit supports for these host systems, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Processor memory

The DS8100 Model 921 offers up to 128 GB of processor memory and the DS8300 Models 922 and 9A2 offer up to 256 GB of processor memory. The Non-Volatile Storage (NVS) scales to the selected processor memory size, which can also help optimize performance.

Data management features

The DS8000 storage unit is designed with the following management features that allow you to securely process and access your data according to your business needs even if it is 24 hours a day and 7 days a week.

RAID

Redundant array of independent disks (RAID) is a method of configuring multiple disk drives in a storage subsystem for high availability and high performance.

The collection of two or more disk drives presents the image of a single disk drive to the system. In the event of a single device failure, data can be read or regenerated from the other disk drives in the array. With RAID implementation, the storage unit offers fault-tolerant data storage. The storage unit supports RAID implementation on the storage unit device adapters. The storage unit supports groups of disk drive modules (DDMs) in both RAID 5 and RAID 10.

RAID 5

RAID 5 is a method of spreading volume data plus data parity across multiple disk drives. RAID 5 increases performance by supporting concurrent accesses to the multiple DDMs within each logical volume.

RAID 10

RAID 10 implementation provides data mirroring from one DDM to another DDM. RAID 10 stripes data across half of the disk drives in the RAID 10 configuration. The other half of the array mirrors the first set of disk drives. In some cases, RAID 10 offers faster data reads and writes than RAID 5 because it does not need to manage parity. However, with half of the DDMs in the group used for data and the other half used to mirror that data, RAID 10 disk groups have less capacity than RAID 5 disk groups.

Arrays across loops

The arrays across loops helps prevent data loss.

The storage unit provides arrays across loops on open-systems, S/390, and zSeries hosts. For RAID 10, arrays across loops provides mirroring across two loops, which prevents loss of the array during loop failure.

Storage System LPARs (logical partitions)

The DS8300 Model 9A2 exploits LPAR technology, allowing you to run two separate storage server images.

Each Storage System LPAR has access to:

- 50 percent of the processors
- 50 percent of the processor memory

- Up to 16 host adapters
- Up to 320 disk drives (up to 96 TB of capacity)

With these separate resources, each Storage System LPAR can run the same or different versions of microcode, and can be used for completely separate production, test, or other unique storage environments within this single physical system. This may enable storage consolidations where separate storage subsystems were previously required, helping to increase management efficiency and cost effectiveness.

Copy Services

Copy Services is a collection of functions that provides disaster recovery, data migration, and data duplication functions. Copy Services runs on the IBM TotalStorage DS8000 storage unit and supports open systems and zSeries environments.

Many design characteristics and advanced functions of the DS8000 contribute to protection of your data. DS8000 has a number of advanced Copy Services functions that are part of the IBM TotalStorage Resiliency family. These functions are supported also on the previous generation of storage systems called the IBM TotalStorage Enterprise Storage Server (ESS).

Copy Services include the following types of functions:

- FlashCopy, which is a point-in-time copy function
- Remote mirror and copy functions (previously known as Peer-to-Peer Remote Copy or PPRC), which includes:
 - IBM TotalStorage Metro Mirror (previously known as Synchronous PPRC)
 - IBM TotalStorage Global Copy (previously known as PPRC Extended Distance)
 - IBM TotalStorage Global Mirror (previously known as Asynchronous PPRC)
- z/OS Global Mirror (previously known as Extended Remote Copy [XRC])

You can manage Copy Services functions through a command-line interface called the IBM TotalStorage DS CLI and a new Web-based interface called the IBM TotalStorage DS Storage Manager. The DS Storage Manager allows you to set up and manage the following types of data-copy functions from any point from which network access is available:

FlashCopy

The FlashCopy feature enables you to create full volume copies of data.

When you set up a FlashCopy operation, a relationship is established between source and target volumes, and a bitmap of the source volume is created. Once this relationship and a bitmap are created, the target volume can be accessed as though all the data had been physically copied. While a relationship between the source and target volume exists, a background process copies the tracks from the source to the target volume.

FlashCopy is an optional function. To use it, you must purchase the point-in-time 2244 function authorization model, which is 2244 Model PTC.

Remote Mirror and Copy

The remote mirror and copy feature is a flexible data mirroring technology that allows replication between volumes on two or more disk storage systems. You can also use this feature for data backup and disaster recovery.

Remote mirror and copy is an optional function. To use it, you must purchase the remote mirror and copy 2244 function authorization model, which is 2244 Model RMC.

DS8000 storage units can participate in remote mirror and copy solutions with the ESS Model 750, ESS Model 800, and DS6000 storage units.

The remote mirror and copy feature can operate in the following modes:

Metro Mirror

Provides real-time mirroring of logical volumes between two DS8000s that can be located up to 300 km from each other. It is a synchronous copy solution where write operations are completed on both copies (local and remote site) before they are considered to be completed.

Global Copy

Copies data nonsynchronously and over longer distances than is possible with Metro Mirror. When operating in Global Copy mode, the source volume sends a periodic, incremental copy of updated tracks to the target volume instead of a constant stream of updates. This causes less impact to application writes for source volumes and less demand for bandwidth resources, while allowing a more flexible use of the available bandwidth.

Global Mirror

Provides a long-distance remote copy feature across two sites using asynchronous technology. Global Mirror operations provide the following benefits:

- Support for virtually unlimited distance between the local and remote sites, with the distance typically limited only by the capabilities of the network and the channel extension technology. This "unlimited" distance enables you to choose your remote site location based on business needs and enables site separation to add protection from localized disasters.
- A consistent and restartable copy of the data at the remote site, created with minimal impact to applications at the local site.
- Data currency where, for many environments, the remote site lags behind the local site on an average of 3 to 5 seconds, minimizing the amount of data exposure in the event of an unplanned outage. The actual lag in data currency that you experience can depend upon a number of factors, including specific workload characteristics and bandwidth between the local and remote sites.
- Dynamic selection of the desired recovery point objective, based upon business requirements and optimization of available bandwidth.
- Session support whereby data consistency at the remote site is internally managed across up to eight DS8000 machines that are located across the local and remote sites.
- Efficient synchronization of the local and remote sites with support for failover and failback modes, helping to reduce the time that is required to switch back to the local site after a planned or unplanned outage.

z/OS Global Mirror

The z/OS Global Mirror function mirrors data on the storage unit to a

remote location for disaster recovery. It protects data consistency across all volumes that you have defined for mirroring. The volumes can reside on several different storage units. The z/OS Global Mirror function can mirror the volumes over several thousand kilometers from the source site to the target recovery site. DS8000 storage complexes support z/OS Global Mirror only on zSeries hosts.

With z/OS Global Mirror, you can suspend or resume service during an outage. You do not have to end your current data-copy session. You can suspend the session, then restart it. Only data that changed during the outage needs to be resynchronized between the copies.

The z/OS Global Mirror function is an optional function. To use it, you must purchase the remote mirror for z/OS 2244 function authorization model, which is 2244 Model RMZ.

z/OS Metro/Global Mirror (3-site z/OS Global Mirror and Metro Mirror)

This mirroring capability uses z/OS Global Mirror to mirror primary site data to a location that is a long distance away and also uses Metro Mirror to mirror primary site data to a location within the metropolitan area. This enables a z/OS 3-site high availability and disaster recovery solution for even greater protection from unplanned outages.

The z/OS Metro/Global Mirror function is an optional function. To use it, you must purchase both of the following functions:

- Remote mirror for z/OS (2244 Model RMZ)
- Remote mirror and copy function (2244 Model RMC) for both the primary and secondary storage units

FlashCopy

The IBM TotalStorage FlashCopy feature provides a point-in-time copy capability for logical volumes. FlashCopy creates a physical point-in-time copy of the data, with minimal interruption to applications, and makes it possible to access immediately both the source and target copies.

The primary objective of FlashCopy is to create a copy of a source volume on the target volume. This copy is called a point-in-time copy. When you initiate a FlashCopy operation, a FlashCopy relationship is created between the source volume and target volume. A FlashCopy relationship is a "mapping" of a FlashCopy source volume and a FlashCopy target volume. This mapping allows a point-in-time copy of the source volume to be copied to the target volume. The FlashCopy relationship exists between the volume pair from the time that you initiate a FlashCopy operation until the DS8000 copies all data from the source volume to the target volume or until you delete the FlashCopy relationship, if it is a persistent FlashCopy.

The point-in-time copy that is created by FlashCopy is typically used when you need a copy of the production data to be produced with minimal application downtime. It can be used for online back up, testing of new applications, or for creating a database for data-mining purposes. The copy looks exactly like the original source volume and is an instantly available, binary copy.

FlashCopy supports the following copy options:

Data Set FlashCopy

Data Set FlashCopy allows a FlashCopy of a data set in a zSeries environment.

Multiple relationship FlashCopy

Multiple relationship FlashCopy allows a source to have FlashCopy relationships with multiple targets simultaneously. This flexibility allows you to establish up to 12 FlashCopy relationships on a given logical unit number (LUN), volume, or data set, without needing to first wait for or cause previous relationships to end.

Refresh target volume (also known as incremental FlashCopy)

Refresh target volume provides the ability to "refresh" a LUN or volume involved in a FlashCopy relationship. When a subsequent FlashCopy operation is initiated, only data that updates the target and the source to the same point-in-time state is copied. The direction of the "refresh" can also be reversed. The LUN or volume that was defined as the target can now become the source for the LUN or the volume that was defined as the source (now the target).

Persistent FlashCopy

Persistent FlashCopy allows the FlashCopy relationship to remain even after the FlashCopy operation completes. You must explicitly delete the relationship.

Establish FlashCopy on existing Remote Mirror and Copy source

The establish FlashCopy on an existing Remote Mirror and Copy source volume option allows you to establish a FlashCopy relationship where the target volume is also the source of an existing remote mirror and copy source volume. This enables you to create full or incremental point-in-time copies at a local site and then use remote mirroring commands to copy the data to the remote site.

This feature is represented by the **Establish target on existing Metro Mirror source** selection in the GUI.

Consistency group commands

Consistency group commands allow the DS8000 to freeze I/O activity to a LUN or volume until you issue the FlashCopy consistency group command. Consistency groups help create a consistent point-in-time copy across multiple LUNs or volumes, and even across multiple DS8000 systems. This function is available through the use of command-line interface commands.

Inband commands over remote mirror link

In a remote mirror environment, inband commands are issued to a source volume of a remote mirror and copy volume pair on a local storage unit and sent across paths (acting as a conduit) to a remote storage unit to enable a FlashCopy pair to be established at the remote site. This eliminates the need for a network connection to the remote site solely for the management of FlashCopy. This function is available through the use of command-line interface commands.

Disaster recovery using Copy Services

One of the main reasons for using Copy Services functions is in preparation for a possible disaster by backing up, copying, and mirroring your data at the production and recovery sites.

You can create a disaster recovery solution using any of the following functions. These functions help improve a disaster recovery process by allowing a much shorter recovery time with little or no data loss.

Note: In a remote mirror and copy environment, when two storage units are set up in two geographically distinct locations, we refer to the sites as Site A (the production site) and Site B (the recovery site).

Failover and failback operations

A failover is the process of temporarily switching production to a backup facility (normally your recovery site) following a planned outage, such as a scheduled maintenance period or an unplanned outage, such as a disaster. A failback operation is the process of returning production to its original location. These operations use remote mirror and copy functions to help reduce the time that is required to synchronize volumes after sites are switched during planned or unplanned outages.

Global Mirror

The Global Mirror function provides a two-site extended distance remote copy option for disaster recovery and backup solution for the zSeries and open systems environments. This solution is based on existing Global Copy and FlashCopy functions. With Global Mirror, the data that the host writes to the storage unit at the production site is asynchronously shadowed to the storage unit at the recovery site. Global Mirror operates over high-speed, Fibre Channel communication links and is designed to maintain a consistent and restartable copy of data at a recovery site that can be located at virtually unlimited distance from the production site.

Comparison of licensed functions

A key decision that you must make in planning for a disaster is deciding which licensed functions to use to best suit your environment.

Table 2 provides a brief summary of the characteristics of the Copy Services features that are available for the storage unit.

Table 2. Comparison of licensed functions

Licensed function	Description	Advantages	Considerations
Metro Mirror	Synchronous data copy at a distance	No data loss, rapid recovery time for distances up to 300 km.	Slight performance impact.
Global Copy	Continuous copy without data consistency	Nearly unlimited distance, suitable for data migration, only limited by network and channel extenders capabilities.	Copy is normally fuzzy but can be made consistent through synchronization.
Global Mirror	Asynchronous copy	Nearly unlimited distance, scalable, and low recovery point objective (RPO). The RPO is the time needed to recover from a disaster; that is, the total system downtime.	RPO might grow when link bandwidth capability is exceeded.

Table 2. Comparison of licensed functions (continued)

Licensed function	Description	Advantages	Considerations
z/OS Global Mirror	Asynchronous copy controlled by z/OS host software	Nearly unlimited distance, highly scalable, and very low RPO.	Additional host server hardware and software is required. The RPO might grow if bandwidth capability is exceeded or host performance might be impacted.

Parallel Access Volumes

The use of parallel access volumes (PAVs) enables a single zSeries server to simultaneously process multiple I/O operations to the same logical volume, which can help to significantly reduce device queue delays. This is achieved by defining multiple addresses per volume.

With dynamic parallel access volumes, the assignment of addresses to volumes is automatically managed to help the workload meet its performance objectives and minimize overall queuing.

You must configure both your storage unit and operating system to use PAVs. You can use the logical configuration definition to define PAV-bases, PAV-aliases, and their relationship in the storage unit hardware. This unit address relationship creates a single logical volume, allowing concurrent I/O operations.

The storage unit supports concurrent or parallel data transfer operations to or from the same volume from the same system or system image for S/390 or zSeries hosts. An S/390 with PAV software support enables multiple users and jobs to simultaneously access a logical volume. Read and write operations can be accessed simultaneously to different domains. (The domain of an I/O operation is the specified extents to which the I/O operation applies.)

DS8000 limitations

The following list describes known limitations for the DS8000.

- The 65,520 cylinder 3390 volume is not supported with z/OS Global Mirror and z/OS Metro/Global Mirror.
- The amount of physical capacity within a 2107 system that can be logically configured for use will be enforced by the 2107 licensed machine code (LMC) to maintain compliance with the extent of IBM authorization established for licensed functions activated on the machine. The 2107 LMC will not allow the logical configuration of physical capacity beyond the extent of IBM authorization (except when activating Standby CoD capacity).
- The deactivation of an activated licensed function, or a lateral change or reduction in the license scope, is a disruptive activity and requires a machine IML (Model 921 and Model 922) or reboot of the affected image (Model 9A2). A lateral change is defined as changing the license scope from fixed block (FB) to count key data (CKD) or from CKD to FB. A reduction is defined as changing the license scope from all physical capacity (ALL) to only FB or only CKD capacity.
- The following activities are disruptive:

- Model conversions. In addition, data may not be preserved during the model conversion.
- Field attachment of the first Model 92E expansion unit to a Model 922.
- Field attachment of the first Model 9AE expansion unit to a Model 9A2.
- Removal of an expansion unit model from the base unit model. Data may not be preserved during this activity.
- Some DS8000 functions are not available or supported in all environments. You can find current information about supported environments, prerequisites, and minimum operating systems levels at the following IBM Web site:
<http://www.ibm.com/servers/storage/disk/ds8000/interop.html>

Planning data migration

The planning and methods of data migration for the DS8000 vary by environment. The DS8000 supports over 90 operating systems. You can migrate data to a storage unit from these host and operating system environments.

When you plan for data migration, consider the following factors:

Note: The following lists do not cover every possibility. They provide a high-level view of some of the tools and factors that you can consider when moving data.

- The system:
 - Is it UNIX based? You will probably use some variation of a logical volume manager.
 - Is it a zSeries or S/390? You will probably use IBM® TotalStorage® Global Mirror, Remote Mirror and Copy (when available).
 - Is it z/OS? You will probably use DFDSS, though there are many choices.
 - Is it VM? You might use DASD Dump Restore or PTAPE.
 - Is it VSE? You might use the VSE fastcopy or ditto commands.

Your system administrator selects the data migration method that is the best compromise between efficiency and impact on the users of the system.
- The storage unit:
 - Are the storage units involved the same with the same level of licensed management code?
 - Are the storage units different? In which case you want to ensure that the new configuration is large enough to accommodate the existing data. You also want to ensure that the virtual disks are similar in configuration to the disk drives that they are replacing.
- Time and complexity involved:
 - Typically data migration requires that updates or changes cease while the movement occurs. Also, depending on the amount of data that you are moving and your migrating method, data could be unavailable for an extended period of time, perhaps several hours.
 - Could the complexity and time involved require the services of IBM through International Global Services? Contact your IBM representative for more information.

How to select a data migration method

Your system administrator selects the data migration method that is the best compromise between efficiency and impact on the users of the system.

Most methods of data migration affect the everyday operation of a computer system. When data is moved, the data must be in a certain state, typically requiring that updates or changes cease while the movement occurs. Depending on the amount of data that you are moving and your migrating method, data could be unavailable for an extended period of time, perhaps several hours. The following factors might contribute to the migration time:

- Creating new logical volumes or file systems
- Modifying configuration files
- Receiving integrity checks

The following items are more than likely among the topics considered by your system administrator to determine the best method to use to migrate your data:

- Management software provides simple robust methods that you can generally use during production without disturbing users.
- The AIX logical volume manager (LVM) provides methods that you can use at any time without disrupting user access to the data. You might notice a small performance degradation, but this is preferable to shutting down databases or requiring users to log off the system.

Notes:

- AIX and HP-UX 11.xx ship with logical volume management (LVM) software as part of the base operating system. LVM provides complete control over all disks and file systems that exist on an AIX system. HP-UX has similar volume management software.
- Sun Microsystems has a basic volume management product called Solstice, which is available for the Solaris systems.
- Linux systems also use the LVM.
- Methods that use backup and restore procedures generally have the most impact on the system usage. They require that databases and file systems be in quiescent states to ensure a valid snapshot of the data.

Chapter 2. Installing the DS8000 Storage Manager interface

The DS8000 Storage Manager is installed using a graphical or silent mode for the Windows and Linux operating systems. It can be accessed using a Web browser from any location that has network access.

You can install the DS8000 Storage Manager on a PC that has the following minimum footprint requirements. Also, ensure you have installed a supported browser.

Requirement	Minimum Value
Disk	600 MB
Memory	512 MB Real Memory
Processor	Pentium 4 Processor at 1.4 GHz

Supported browsers:

- Internet Explorer 6.x
- Netscape 6.2
- Netscape 7.x

Supported operating systems:

- Windows Server 2003 Enterprise and Standard Edition
- Windows 2000 Advance Server, Server, and Professional, each with SP4
- Windows XP Professional SP1 and SP1a
- Windows XP Professional SP2
- Linux (Red Hat Enterprise Linux 2.1)

Note: The DS8000 Storage Manager is not supported on any Windows 64-bit operating system.

Note: Animations must be turned on in your browser if you want to observe the installation progress bars that are associated with the DS Storage Manager installation. Use the following instructions for your specific browser:

- Internet Explorer
 1. From the **Tools** menu, select **Internet Options**.
 2. Select the **Advanced** tab and scroll down to the **Multimedia** section.
 3. Check **Play animation in web pages**.
- Netscape
 1. From the **Edit** menu, select **Preferences**.
 2. Double-click on **Privacy and Security**.
 3. Select **Images** and select **as many times as the image specifies** in the **Animated image should loop** section.

Firewall installation alert

Note: During the installation process, it is possible to receive several alerts for Java processes that are trying to access ports on the local computer. When this occurs, allow the Java processes to access the ports; otherwise, the installation might fail.

After you have installed the DS8000 Storage Manager, the following results occur:

- Activation of the IBM TotalStorage DS Storage Manager server and the IBM TotalStorage DS Network Interface server. These servers are set to automatic startup so that when you start your computer these servers are automatically activated.
- Activation of the IBM TotalStorage DS8000 Storage Manager application, which includes the real-time and simulated manager components. Both these components can be installed on the same machine and are integrated into the user interface. They are designed to help you create and manage the physical and logical configurations of your storage complexes, storage units, and storage images. Plus, the real-time manager component provides you the opportunity to use the Copy Services features you have purchased.

Real-time manager

This component (preinstalled on the HMC) provides logical configuration and Copy Services features to a network attached DS8000.

Simulated manager

This component (installed using the IBM TotalStorage DS8000 Storage Manager installation CD) provides the ability to create or modify logical configurations when you are disconnected from the network. After creating the configuration, you can save it and then apply it to a network-attached DS8000 at a later time.

Installing the IBM TotalStorage DS Storage Manager on the Windows operating system using the graphical mode

Use the following steps to install the IBM TotalStorage DS Storage Manager in your Windows environment using the graphical mode. When you are installing on a DS8000 storage unit, you are generally installing the offline management type. This management type allows you to plan and construct storage unit configurations without being connected to the network. However, you cannot use the Copy Services or parallel access volumes (PAVs) features while disconnected from the network.

Before you install the IBM TotalStorage DS Storage Manager, verify that the prerequisite software and hardware are installed on your system. The installation program checks for prerequisites and stops if any prerequisites are missing.

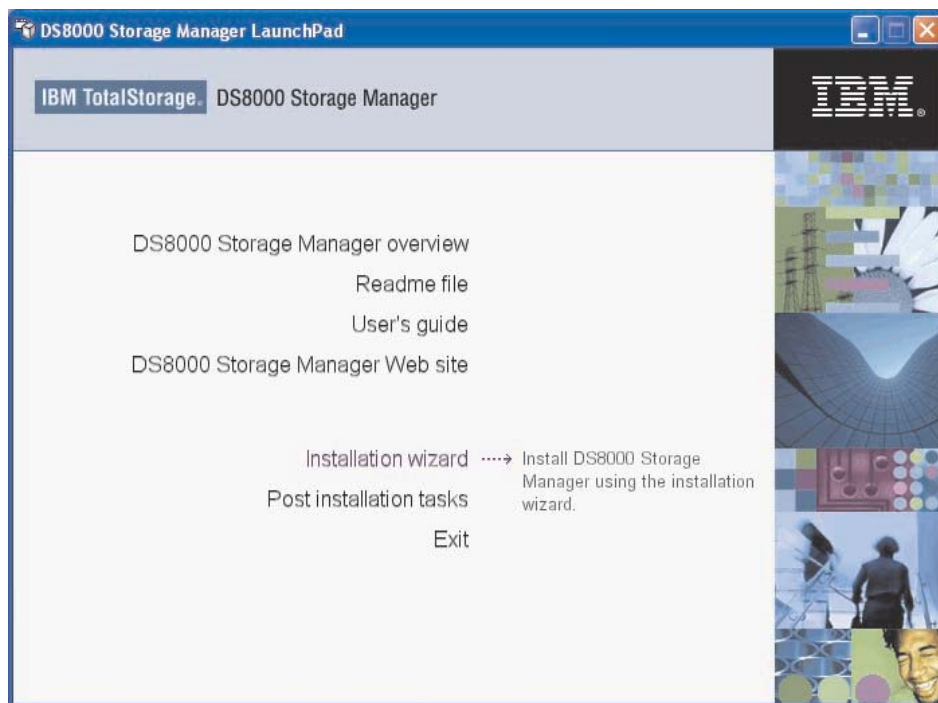
Note: The DS Storage Manager is not supported on any Windows 64-bit operating system.

1. Log on as a user with administrator authority.
2. Insert the IBM TotalStorage DS Storage Manager CD into the CD-ROM drive. The IBM TotalStorage DS Storage Manager program starts within 15 - 30 seconds if you have autorun mode set on your system. The LaunchPad window is displayed.

If the LaunchPad window does not display, go to the CD-ROM drive using Windows Explorer or a command prompt and perform one of the following steps:

- a. Type *LaunchPad* at the command prompt and press **Enter**. The LaunchPad window is displayed.
- b. Locate and double-click the **LaunchPad.bat** reference in Windows Explorer.

Note: If you are viewing the folder with Windows Explorer with the option selected to hide the extensions for unknown file types, find the LaunchPad file with the file type of MS-DOS Batch file.



Launchpad window

3. Choose one of the following options that are listed on the LaunchPad window:

DS8000 Storage Manager overview

Provides information about the IBM TotalStorage DS Storage Manager software.

Readme file (recommended selection)

Provides last minute product information that was not provided in these installation instructions.

User's guide

Provides specific installation instructions.

DS8000 Storage Manager Web site

Provides information from the product Web site.

Installation wizard

Starts the IBM TotalStorage DS Storage Manager installation program.

Post installation tasks

Provides information about configuring the IBM TotalStorage DS Storage Manager.

Exit

Exits the IBM TotalStorage DS Storage Manager LaunchPad program.

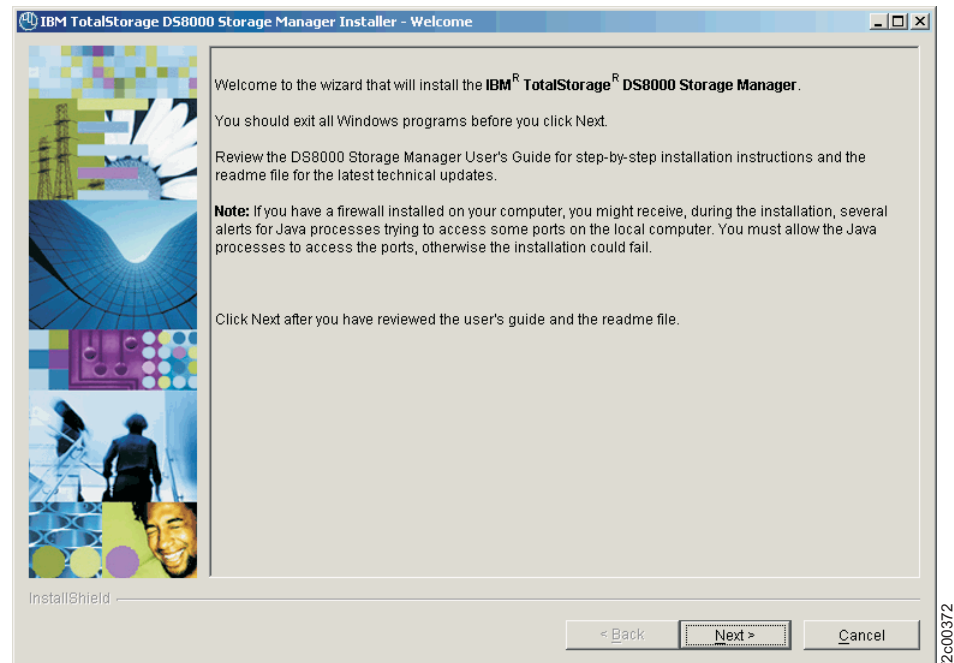
4. Click the **Readme file** selection on the LaunchPad to check for information that might supersede the information in this guide.
5. Click the **Installation wizard** selection on the LaunchPad to start the installation program.

Note: The LaunchPad window remains open behind the installation wizard so that you can access product information during the installation process.

There might be a slight delay while the software loads on your system. After the software loads, a DOS prompt window opens to display the following message:

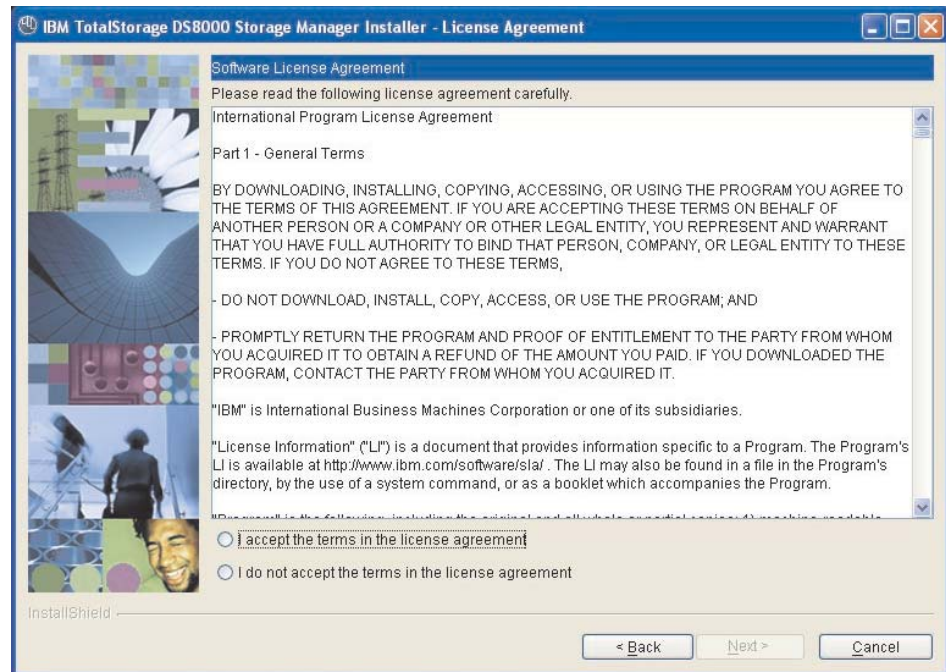
```
Initializing InstallShield Wizard...
Preparing Java (tm) Virtual Machine .....
.....
```

The Welcome window of the IBM TotalStorage DS Storage Manager installation program displays if no problems are discovered during the initial system check. If an error is discovered (for example, the operating system does not match the prerequisite), an error message is displayed and the installation program exits.



DS8000 Storage Manager Installer Welcome window

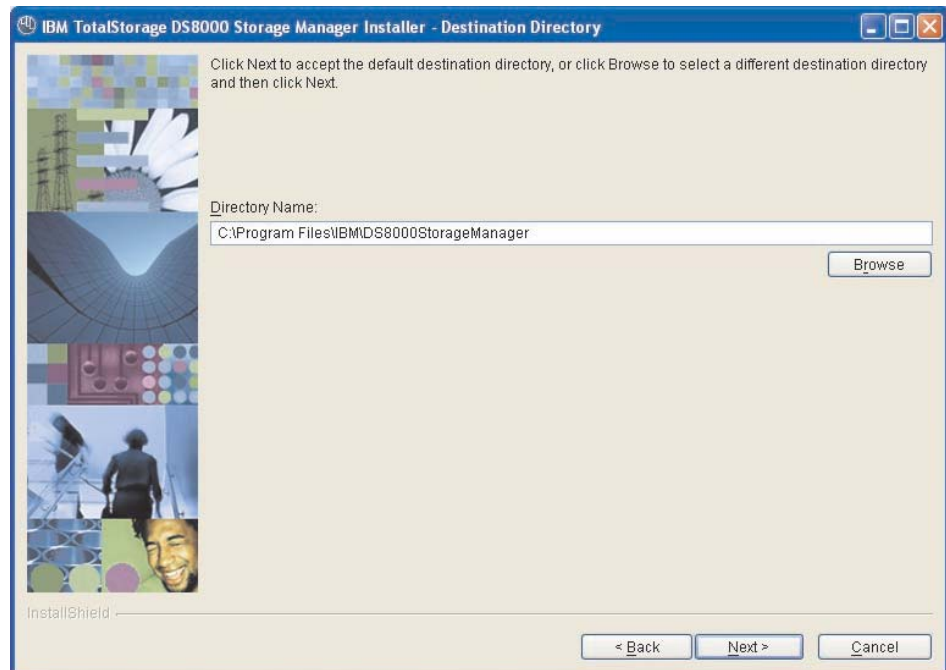
6. Click **Next** to continue, or click **Cancel** to exit the installation. When you click **Next**, the License Agreement window displays.



License Agreement window

Note: Anytime you click **Cancel** on any of the installation windows, a message asking for you to confirm that you want to exit is displayed.

7. Read the license agreement and click your acceptance or nonacceptance of the agreement. If you accept, the **Next** button is highlighted. Click **Next** to continue or click **Cancel** to exit the installation. When you click **Next**, the Destination Directory window is displayed.



Destination Directory window

8. Choose the directory on the Destination Directory window where you want to install the application by using one of the following options. Clicking **Next** after

any one of the options causes the server installation checking window (a window for each server) to display if the DS Storage Manager is installed on your system, or if the Server Parameters window is displayed.

- a. Click the **Next** button to accept the default (recommended) directory that is shown in the window.
- b. Type a fully qualified directory path into the Directory Name path to designate your own directory. Then click the **Next** button.
- c. Click the **Browse** button and a directory window is displayed. Scroll to the directory that you want to use for the application. Click it and click **OK**. The directory path is displayed in the Directory Name field. Click **Next** to continue.
- d. Click **Cancel** to exit the installation process.

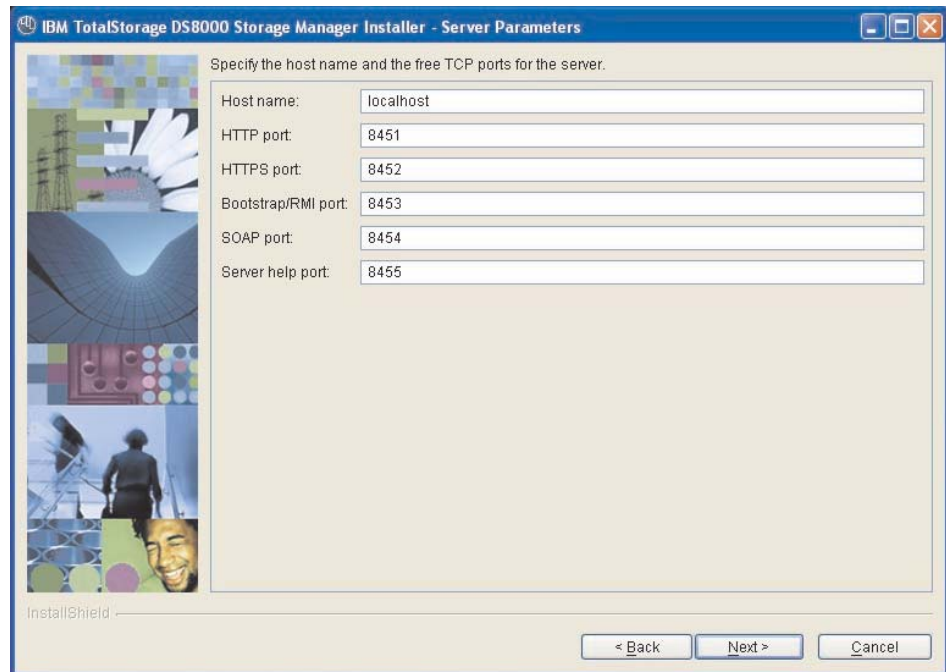
9. Observe the message on each of the next two Installation Checking windows and click **Next** on each window to continue. If the servers have not been installed, the Server Parameters window is displayed with default values in the input fields. Go to Step 10 to proceed with the Server Parameters window.

The DS Storage Manager Server Installation Checking window is displayed if the DS Storage Manager server is already installed on your system. A message is displayed on the window that indicates which version of the server is installed and whether the installation process will install a newer version of the server. To continue the installation process, you must click **Next** or click **Cancel** to exit the installation process.

When you click **Next**, the DS Network Interface Server Installation Checking window is displayed with a message. This message indicates the version of the DS Network Interface server that is installed on your system and whether the installation process will install a newer version of the server. To continue the installation process, you must click **Next** or click **Cancel** to exit the installation process.

When you click **Next**, the Installation Confirmation window is displayed. If your installation process has taken you through these two checking windows, go to Step 16 on page 34.

10. Click **Next** to accept the default values that are displayed on the Server Parameters window. All fields must contain a value before you can proceed to the next step.



Windows Server Parameters window

You can change the default values by supplying the following information:

- a. **Host Name** — The host name of the physical machine where the DS Storage Manager Server is installed. The host name must resolve to a physical network node on the server. When multiple network cards exist in the server, the host name or IP address must resolve to one of the network cards. The value that you specify for the host name is used as the value of the `hostName` property in WebSphere Application Server configuration files. Specify the host name value in one of the following formats:
 - The fully qualified domain name servers (DNS) host name string; for example, *xmachine.manhattan.ibm.com*.
 - The default short DNS host name string; for example, *xmachine*.
 - The numeric IP address; for example, *9.127.255.3*.
 - The generic name *localhost* or its equivalent IP address; for example, *127.0.0.1*.

Note: The specified value for the host name is stored by WebSphere Application Server into its internal configuration files. If you specified the host name in the fully qualified form or the short host name string, do not change the host name after the DS Storage Manager Server installation. If you do, the host name and the name that you specified for the configuration file no longer match. If you do change the host name, you must manually edit all the WebSphere Application Server configuration files that contain references to the old name to reflect the new name. If you specified the host name in the form of the numeric IP address, you must use a fixed IP address for the computer (you cannot use DHCP to obtain a dynamic IP address for it). The above restrictions do not exist if you specified the host name in the form as *localhost*.

- b. **HTTP Port** — The port number that the DS Storage Manager Server uses. The port number must be a port that is not being used by another process on the system. After the DS Storage Manager Server is installed, you must

include this port number in the URL to access the console. That URL is the protocol name, plus the fully qualified host name, plus the port, plus DS8000/console; for example, `http://myhost.com:8451/DS8000/console`.

- c. **HTTPS Port** — The port that the DS Storage Manager Server uses for secure HTTP transport (HTTPS).
- d. **Bootstrap/RMI Port** — The port that the DS Storage Manager Server uses for the bootstrap port.
- e. **SOAP Port** — The port that the DS Storage Manager Server uses for Simple Object Access Protocol (SOAP).
- f. **Server help Port** — The port for the Eclipse help Server. This port is the one that the help system (based on Eclipse technology) uses to receive requests for help files.

Note: The values that are used in the Server Parameters window fields must not conflict with existing port assignments that are already on the system or reserved for use by another application.

11. Click **Next** after you have ensured that all the input fields on the Server Parameters window are complete. The SSL Configuration window is displayed. This window allows you to supply the fully qualified name of the two server key files generated before or during the DS8000 Storage Manager installation.

IBM TotalStorage DS8000 Storage Manager Installer - SSL Configuration

Please enter the following information required to set SSL configuration for IBM TotalStorage DS8000 Storage Manager.
 'Key file name' stands for the fully qualified path to the key file.
 'Key file password' stands for the password for accessing the key file.
 'Trust file name' stands for the fully qualified path to a trust file.
 'Trust file password' stands for the password for accessing the trust file.

Select one of the following options concerning the SSL certificates:

☒ Generate the self-signed certificates during installation, or
☐ Use existing SSL certificates

Key file name:

Key file password:

Confirm key file password:

Trust file name:

Trust file password:

Confirm trust file password:

InstallShield

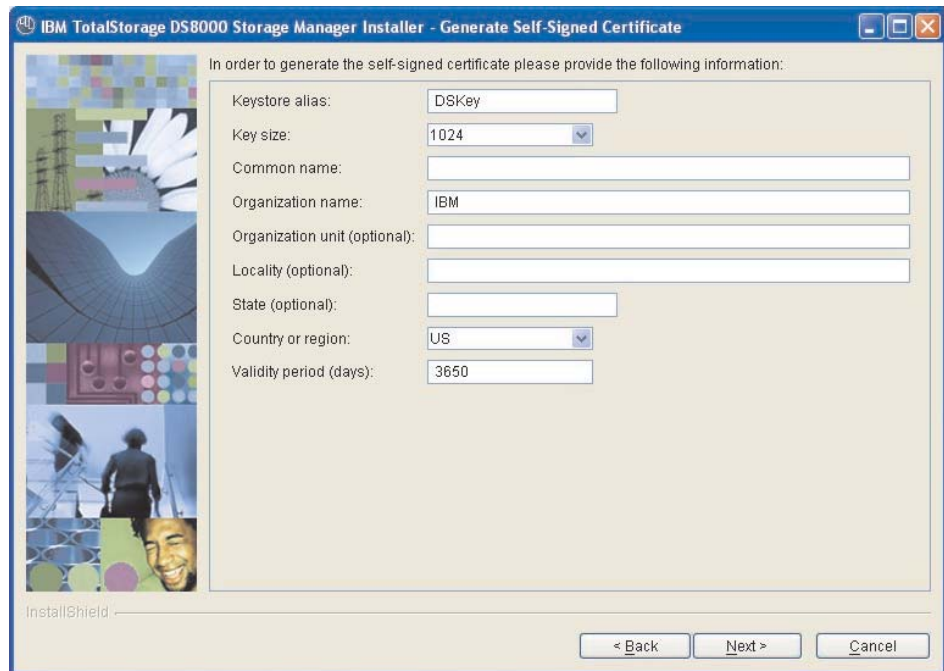
< Back Next > Cancel

Windows SSL Configuration window

12. Follow the instructions on the window.
 - a. First, designate whether you want to generate a new SSL certificate by selecting **Generate the self-signed certificates during installation**. This is the most likely selection when you are installing DS8000 Storage Manager for the first time. You can also select **Use existing SSL certificates**. This choice is most likely made when the key files already exist and you do not want to generate new key files.
 - b. Complete the rest of the fields on the window. Provide the fully qualified path to the key file and trust file by using the Browse button. The key file and trust file are identified by an extension of .jks.

- c. Supply a password for each file. The password must be a minimum of six-characters and it cannot contain any blanks.
13. Click **Next** to continue the installation process. If you have selected **Generate the self-signed certificates during installation**, the Generate Self-Signed Certificate window is displayed. If you selected **Use existing SSL certificates**, the Installation Confirmation window is displayed.

Note: When you select **Generate self-signed certificates during installation** and the key files already exist on the system at the specified location, clicking **Next** causes a message to display. It asks you to designate whether you want to have the system delete existing application server certificates if they exist on the location and to generate others. The default is **NO** which means keep the existing application server certificates.



Windows Generate Self-Signed Certificate window

14. Complete each of the input fields with the required information. Use the following as a guide to provide the needed information.

Keystore alias

This is the alias for the self-signed certificate. It uniquely identifies the certificate within the keystore file. It is a good practice to use a unique name related to the server name.

Key size

Two key sizes are available, 512 and 1024. 1024 is the default.

Common name

This name is the primary, universal identity for the certificate. If your system contains the secured WebSphere environment, this name must be valid in the configured user registry.

Organization name

The name of your company

Organization unit (optional)

The name of a department or division within your company. Remember

that you are building a profile for the certificate. As a security measure the more specific the information the more secure your system.

Locality (optional)

The city or location where your company resides.

State (optional)

The state or province where your company resides.

Country or region

The two-character designator that identifies the country where your company is located. Use the pull-down menu to make this selection.

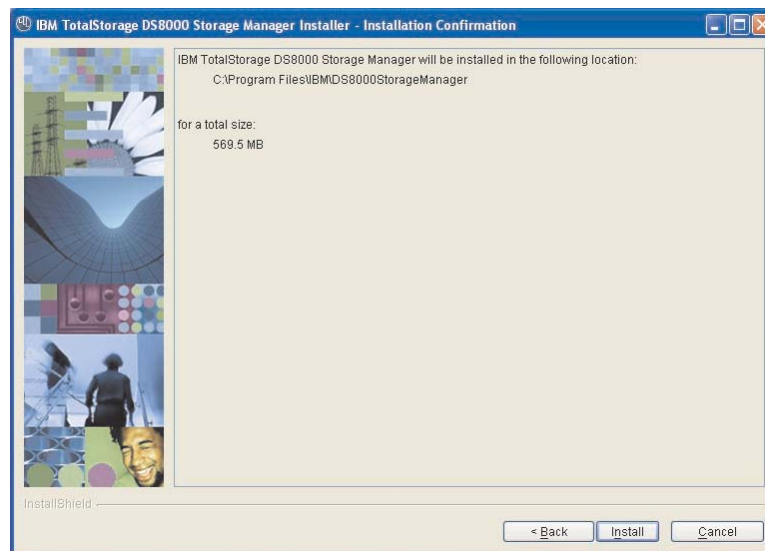
Validity period (days)

Specifies the lifetime of the certificate in days.

The information that you provide is used to build a profile for the certificate. This information is used during identity processing to ensure validation of any processing that is occurring. It is part of the security measures that are used during SSL connectivity.

Click **Next** to continue, or click **Cancel** to exit the installation.

15. Click **Next**. The Installation Confirmation window is displayed. This window displays both the location where the product will be installed and the total size needed for the installation.



Windows Installation Confirmation window

16. Click the **Install** button on the Installation Confirmation window to begin the installation process. There are several progress windows that are displayed. There is no required interaction on your part for each of the progress windows that are displayed. However, you can choose to cancel (not recommended) the installation on any of the progress windows with varying consequences.

The installation process performs the following actions:

- a. If the two servers (DS Storage Manager Server and DS Network Interface Server) are already installed on your system, they are stopped in the following order of windows:
 - 1) The Embedded IBM WebSphere Application Server - Express server (part of the DS Storage Manager Server) window is stopped first.
 - 2) The service window (DS Network Interface Server) is stopped next.

3) The WS Help System (part of the DS Storage Manager Server) window is stopped, if it was not stopped before by the Embedded IBM WebSphere Application Server.

b. If one or both of the servers are not installed or have to be upgraded on your system, they are installed or upgraded in the following order (the progress of the installation is indicated on the associated progress window):

- 1) DS Storage Manager Server Installation Progress window
- 2) DS Network Interface Server Installation Progress window

Note: You can click **Cancel** (not recommended) during the DS Network Interface Server installation process. The process does not stop immediately when the **Cancel** button is clicked. Rather, the process continues to install all the files that are associated with this part of the installation. These files remain installed and are not reinstalled upon reactivation of the installation process.

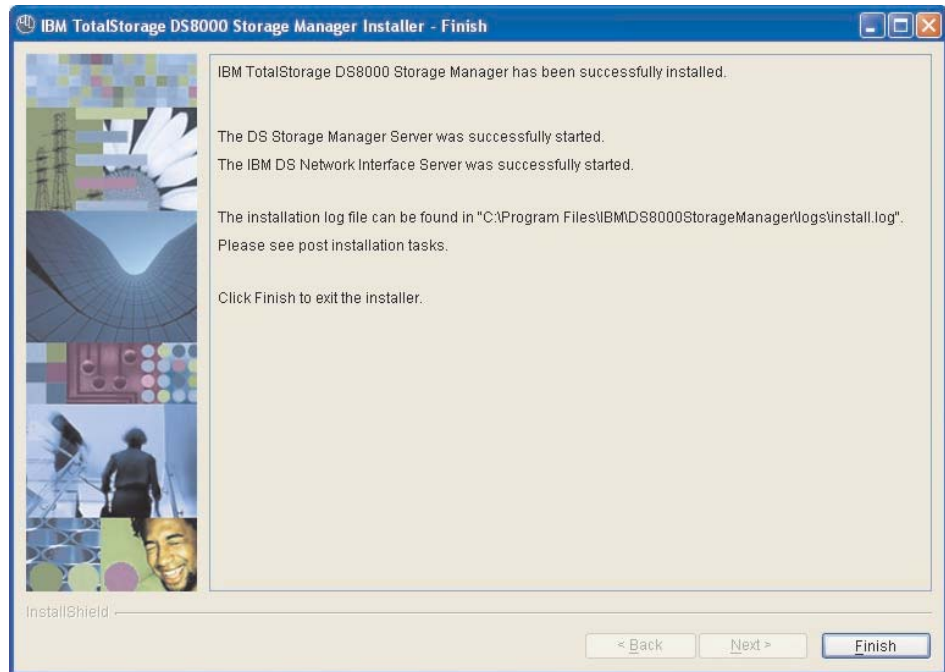
A confirmation message is displayed that asks you to confirm that you want to cancel the installation.

c. The Components Installation Progress (displaying the installation or upgrade progress of the DS Storage Manager product applications) is displayed after the servers have been installed or upgraded.

Note: You can click **Cancel** (not recommended) during the components installation process. The installation process stops immediately when the **Cancel** button is clicked. A window with a confirmation message is displayed. For a new installation, when you confirm that you want to stop the process, all the files that have been copied up to the point that the **Cancel** button was clicked are uninstalled. You then are exited from the installation process.

When this part of the installation is completed, the system starts both servers: first the DS Network Interface Server and then the DS Storage Manager Server. Wait for the servers to be started before going to the next step. When the servers have been started, the Installer Finish window is displayed.

17. Click **Finish** to exit the installation process. When the installation process exits, a file (postinstallation.txt) is displayed, containing the postinstallation tasks. Follow the instructions in this text file to complete the postinstallation tasks.



Windows DS8000 Storage Manager Installer Finish window

Note: If the installation fails, you must exit the installation process and check the install log for error messages.

18. If the Finish window indicates that the installation fails, check the installation log for error messages. The installation log is located in `xxx\logs\install.log`, where `xxx` is the destination directory where the IBM TotalStorage DS Storage Manager is installed (for example, `c:\Program Files\IBM\DS8000 Storage Manager`) .

Errors fall into two categories: system or incorrect values that are designated during the installation process. Use the following guidelines to correct these errors.

- If the error is due to a system problem, correct it and reinstall the DS8000 Storage manager using either the interactive or silent mode of installation.
 - If the error is due to a wrong installation parameter value, restart the installation using the steps that are described in this procedure or the steps that are described in the silent mode of installation. Navigate to the Server Parameters window and insert the correct values. Then finish the installation process.
19. Complete the postinstallation tasks. If, when exiting the installation wizard, the `postinstallation.txt` file does not automatically open, manually open it from the LaunchPad window, and follow the instructions to complete the postinstallation tasks.
 20. Exit the LaunchPad program by clicking **Exit** on the LaunchPad window.

Installing the IBM TotalStorage DS Storage Manager on the Windows operating system in unattended (silent) mode

Use the following steps to install the IBM TotalStorage DS Storage Manager in your Windows environment using the unattended (silent) mode. When you are installing on a DS8000 storage unit, you are generally installing the offline management type. This management type allows you to plan and construct storage unit configurations without being connected to the network. However, you cannot use the Copy Services or parallel access volumes (PAVs) features while disconnected from the network.

Before you install the IBM TotalStorage DS Storage Manager, verify that the prerequisite software and hardware are installed on your system. The installation program checks for prerequisites and stops if any prerequisites are missing.

Note: The DS Storage Manager is not supported on any Windows 64-bit operating system.

The unattended (silent mode) installation option allows you to run the installation program unattended. Use this method of installation to customize a response file and to issue a command from a command prompt window. The response file is a template on the IBM TotalStorage DS Storage Manager CD.

Perform the following steps to install the IBM TotalStorage DS Storage Manager in your Windows environment using the unattended mode:

1. Log on to your Windows system as an administrator.
2. Insert the IBM TotalStorage DS Storage Manager installation CD into the CD-ROM drive. If the interactive mode starts, click the **Exit** selection on the LaunchPad window to exit.
3. Locate the response file (responsefile.txt) in the root directory of your CD and copy it to a directory of your choosing on your system, or open a command prompt window and generate the template response file.

To generate the template response file, perform the following steps:

- a. Ensure that your command prompt is located at the root directory of the CD-ROM drive.
 - b. Type `setup.exe -options-template <responsefile-path>\<responsefile>`, where *responsefile-path* is the path where you want the response file to be created, and *responsefile* is the name of the response file that you want to create.
 - c. Press the **Enter** key and wait for the program to create the response file.
4. Open the response file using a text editor and modify the default options with the values that you want.

Follow these guidelines when you modify the default options:

- Remove the three # characters from the beginning of each line that contains a parameter (if you generated the response file as shown before). Change the parameter default value to the value that you want for that option. You *must* enclose all values in double quotation marks ("*<value>*").
- Verify that all paths generated or modified in the response file are correct. There are some cases when the generated file displays an incorrect string representing the path of the installation location. For example, the parameter `"-P installLocation = "C:Program FilesMy Product"` is wrong and must be `"-P product.installLocation = "<directory where you want the DS Storage Manager installed>"`

- Ensure that the entered values for the ports values are available and not used by other applications on your system. Use the command "*netstat -a*" to verify the ports in use on your system.
- Specify two passwords for the SSL keys files. These passwords must observe the following criteria:
 - The passwords must contain at least 6 characters.
 - Use the following options within the responsefile.txt file to set these options:

- For the keystore use:

```
-W wasSslConfiguration.keystorePassword="your_keystore_password"
```

- For the truststore use:

```
-W wasSslConfiguration.truststorePassword="<your_truststore_password>"
```

Note: These options do not have default values. If you do not set these values, the installation process fails.

5. Save the modifications to the response file.
6. Type the following command at the command prompt and press the **Enter** key on your keyboard to start the installation process in silent mode: *setup.exe -options <responsefile-path>\<responsefile> -silent* where
 - *responsefile-path* — represents the path where the response file resides.
 - *responsefile* — represents the relative name of the response file that you used or created (for example, responsefile.rsp or responsefile.txt).
7. Wait for the installation program to install the product. This can take 5 - 10 minutes.

Note: During this processing, the installation process checks to see if the DS Storage Manager server and DS Network Interface Server are already installed and if they are the current version or have to be upgraded. If they are already installed and need to be upgraded, the installation process performs this function using the already set values from the previous server installation, found on the system in the server configuration files, for all the associated values. The specified values in the response file are ignored.

8. Check the install log file for any possible error messages. This file is located in the **xxx\logs\install.log** directory, where xxx is the destination directory where the IBM TotalStorage DS Storage Manager is installed.
9. Start the IBM TotalStorage DS Storage Manager, if no errors are evident.
10. Perform the postinstallation tasks when the DS Storage Manager has been installed successfully. You can access the instructions for the postinstallation tasks from the postinstallation.txt file, in the doc directory on the product CD.

The following is an example of the template response file that is created when you process the previous steps.


```
#####
#
# InstallShield Options File
#
# Wizard name: Setup
# Wizard source: setup.jar
# Created on: Tue Jul 27 09:02:47 EEST 2004
# Created by: InstallShield Options File Generator
#
# This file contains values that were specified during a recent execution of
# Setup. It can be used to configure Setup with the options specified below when
# the wizard is run with the "-options" command line option. Read each setting's
# documentation for information on how to change its value.
#
# A common use of an options file is to run the wizard in silent mode. This lets
# the options file author specify wizard settings without having to run the
# wizard in graphical or console mode. To use this options file for silent mode
# execution, use the following command line arguments when running the wizard:
#
# -options "responsefile.txt"
#
#####

-silent

#####
#
# IBM TotalStorage DS8000 Storage Manager Install Location
#
# The install location of the product. Specify a valid directory into which the
# product should be installed. If the directory contains spaces, enclose it in
# double-quotes.
#

-P product.installLocation="C:\Program Files\IBM\DS8000StorageManager"

#####
#
# User Input Field - hostname
#
# The fully-qualified host name of the machine where DS Storage Manager Server
# will be installed.
#

-W wasExpressConfig.hostname="localhost"

#####
#
# User Input Field - httpPort
#
# The port number that the HTTP transport in the application server will use.
# This value must not conflict with existing port assignments on the system.
#

-W wasExpressConfig.httpPort="8451"

#####
#
# User Input Field - httpsPort
#
# The port number that the HTTPS transport in the application server will use
# for secure HTTP transport. This value must not conflict with existing port
# assignments on the system.
#

-W wasExpressConfig.httpsPort="8452"
```

```
#####
#
# User Input Field - bootstrapPort
#
# The address for the bootstrap function and the port number for the Java Remote
# Method Invocation (RMI) connector in the application server. This value must
# not conflict with existing port assignments on the system.
#

-W wasExpressConfig.bootstrapPort="8453"

#####
#
# User Input Field - soapPort
#
# The address for the Simple Object Access Protocol (SOAP) connector in the
# application server. This value must not conflict with existing port
# assignments on the system.
#

-W wasExpressConfig.soapPort="8454"

#####
#
# User Input Field - helpPort
#
# The port for Eclipse
# This value must not conflict with existing port assignments on the system.
#

-W wasExpressConfig.helpPort="8455"

#####
#
# User Input Field - certificateOption
#
# Select one of the following options concerning the SSL certificates.
# Legal values are: "generate" (the default value) - Generate the
# self-signed certificates during installation. "use" - Use existing
# certificates. For example, to specify that the "generate" option is selected,
# use -W wasSslConfiguration.certificateOption="generate"
#

-W wasSslConfiguration.certificateOption="generate"

#####
#
# User Input Field - keystoreFileName
#
# The absolute path of the keystore file.
#

-W wasSslConfiguration.keystoreFileName="C:\Program
Files\IBM\SMServer\keys\SMServerKeyFile.jks"

#####
#
# User Input Field - keystorePassword
#
# The password for the keystore file.
#

-W wasSslConfiguration.keystorePassword="<your_keystore_password>"
```

```
#####
#
# User Input Field - truststoreFileName
#
# The absolute path of the truststore file.
#

-W wasSslConfiguration.truststoreFileName="C:\Program
Files\IBM\SMServer\keys\SMServerTrustFile.jks"

#####
#
# User Input Field - truststorePassword
#
# The password for the truststore file.
#

-W wasSslConfiguration.truststorePassword="<your_truststore_password>"

#####
#
# Delete server certificates option
#
# Option to delete certificates files if they exist. Legal values are: "yes" - in
# order to delete application server certificates if they exist on the location,
# and to generate others, or "no" (the default value) - in order to keep the
# existing application server certificates.
#

-G deleteCertificates=yes

#####
#
# User Input Field - keystoreAlias
#
# The alias for the self-signed digital certificate, which is used to uniquely
# identify the certificate within the keystore file. If you have only one
# certificate in each keystore file, you can assign any value to the label.
# However, it is good practice to use a unique label related to the server name.
#

-W CertificateParams.keystoreAlias="DSKey"

#####
#
# User Input Field - keySize
#
# The key size. That must be between 512 and 1024 and must be multiple of
# 64b. The only two allowed values are 512 or 1024. If you enter other values it
# takes the default value 1024.
#

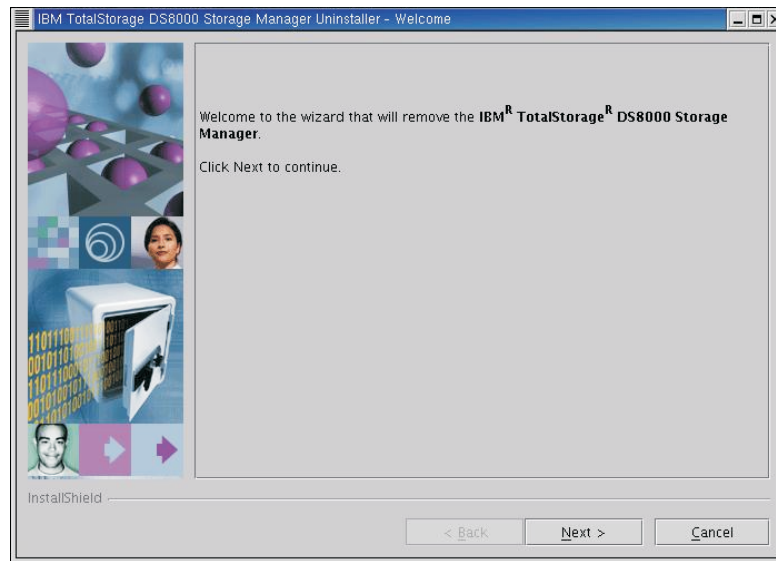
-W CertificateParams.keySize="1024"

#####
#
# User Input Field - commonName
#
# The common name is the primary, universal identity for the certificate; it
# should uniquely identify the principal that it represents. In a WebSphere
# environment, certificates frequently represent server principals, and the
# common convention is to use common names of the form "host_name" and
# "server_name". The common name must be valid in the configured user registry
# for the secured WebSphere environment.
#
```

Removing the IBM TotalStorage DS Storage Manager from a Windows operating system

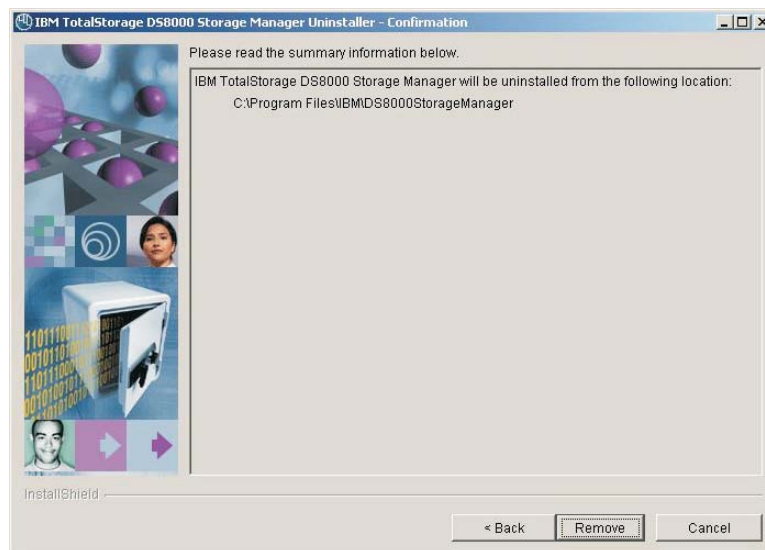
You can remove the IBM TotalStorage DS Storage Manager from your Windows operating system by using the Windows Add/Remove Programs facility.

1. Navigate to your Control Panel and open the Add/Remove program facility.
2. Scroll the list of currently installed programs and click the listing for IBM TotalStorage DS8000 Storage Manager.
3. Click the **Change/Remove** button, and the Welcome window for the Uninstaller is displayed.



Uninstaller Welcome window

4. Click **Next** to continue or click **Cancel** to exit the removal process. When you click **Next**, the Confirmation window is displayed showing the directory from which the IBM TotalStorage DS Storage Manager GUI program will be removed.



Uninstaller Windows Confirmation window

5. Click **Remove** to continue or **Cancel** to stop the removal and exit the uninstallation process. Click **Back** to return to the previous window. When you click **Remove**, the Uninstallation Progress window is displayed. When the uninstallation process is finished, the Finish window is displayed, which contains a statement about the success or failure of the uninstallation process.

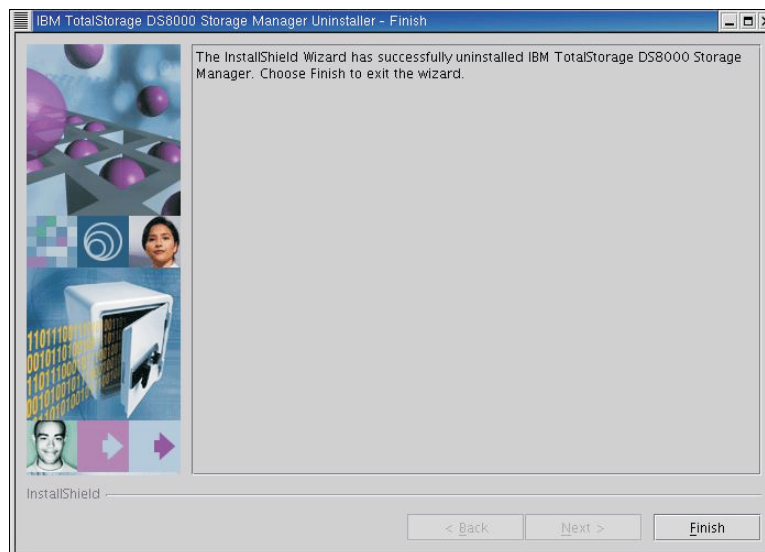
During the uninstallation process, the following system checks occur:

- A check of your system to detect if DS Storage Manager Server is used by other programs on the system and if not used, removes it from the system. Otherwise, the DS Storage Manager Server is not removed from your system.
- A check of your system to detect if DS Network Interface Server is used by other programs on the system and if not used, removes it from the system. Otherwise, the DS Network Interface Server is not removed from your system.
- A check to determine whether any of the files to be removed are locked because they are in use by another process. If a locked file has been detected, you can perform the following:
 - a. From Windows services, set the **IBM WebSphere Application Server V5 - DS Storage Manager** and **IBM DS Network Interface Server** services to **Manual** startup type.
 - b. Reboot the computer.
 - c. Proceed once again with the removal of the DS Storage Manager.

Another reason files could be locked is because you changed the host name you installed the DS Storage Manager. See the DS Storage Manager installation instructions for more information on host names and how to resolve a changed host name.

When the uninstallation process is finished, the Uninstaller Finish window is displayed. The Uninstaller Finish window provides a statement about the success or failure of the uninstallation process.

If the uninstall program cannot remove some information from your system, the Reboot window is displayed. You must reboot so that previously locked files are released and automatically deleted.



Uninstaller Finish window

6. Click **Finish** to complete the removal process and exit the GUI part of the uninstallation program.

7. Close the Add/Remove Programs window.
8. Restart your system, if required (now or later) to complete the removal process.

Removing the DS8000 Storage Manager from a Windows operating system using the silent mode

You can remove the IBM TotalStorage DS Storage Manager from your Windows operating system by using the silent (unattended mode).

1. Log on to your Windows operating system with administrator authority.
2. Open the Command Prompt window.
3. Navigate to the directory where the DS8000 Storage Manager is installed.
4. Type `<dest_path>_uninst\uninstall.exe -silent`, where `<dest_path>` is the path where the DS8000 Storage Manager is installed.
5. Wait for the program to remove the DS8000 Storage Manager.
6. Restart your system if required (now or later) to complete the removal of the program.

Upgrading the DS Storage Manager on a Windows operating system using the graphical mode

Use this process to upgrade the DS Storage Manager to the latest level using the graphical mode for a Windows operating system.

You must have the latest DS Storage Manager CD-ROM or you must download the latest version from the Web site.

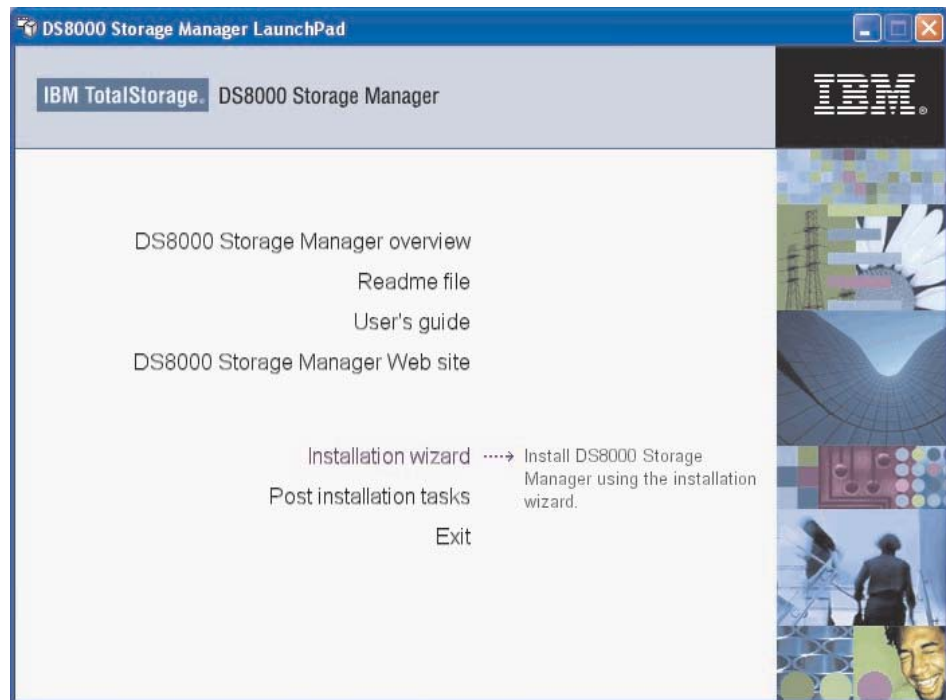
Note: The DS Storage Manager is not supported on any Windows 64-bit operating system.

1. Log on as a user with administrator authority.
2. Insert the IBM TotalStorage DS Storage Manager CD into the CD-ROM drive. The IBM TotalStorage DS Storage Manager program starts within 15 - 30 seconds if you have autorun mode set on your system. The LaunchPad window is displayed.

If the LaunchPad window does not display, go to the CD-ROM drive using Windows Explorer or a command prompt and perform one of the following steps:

- a. Type *LaunchPad* at the command prompt and press **Enter**. The LaunchPad window is displayed.
- b. Locate and double-click the **LaunchPad.bat** reference in Windows Explorer.

Note: If you are viewing the folder with Windows Explorer with the option selected to hide the extensions for unknown file types, find the LaunchPad file with the file type of MS-DOS Batch file.



Launchpad window

3. Choose one of the following options that are listed on the LaunchPad window:

DS8000 Storage Manager overview

Provides information about the IBM TotalStorage DS Storage Manager software.

Readme file (recommended selection)

Provides last minute product information that was not provided in these installation instructions.

User's guide

Provides specific installation instructions.

DS8000 Storage Manager Web site

Provides information from the product Web site.

Installation wizard

Starts the IBM TotalStorage DS Storage Manager installation program.

Post installation tasks

Provides information about configuring the IBM TotalStorage DS Storage Manager.

Exit

Exits the IBM TotalStorage DS Storage Manager LaunchPad program.

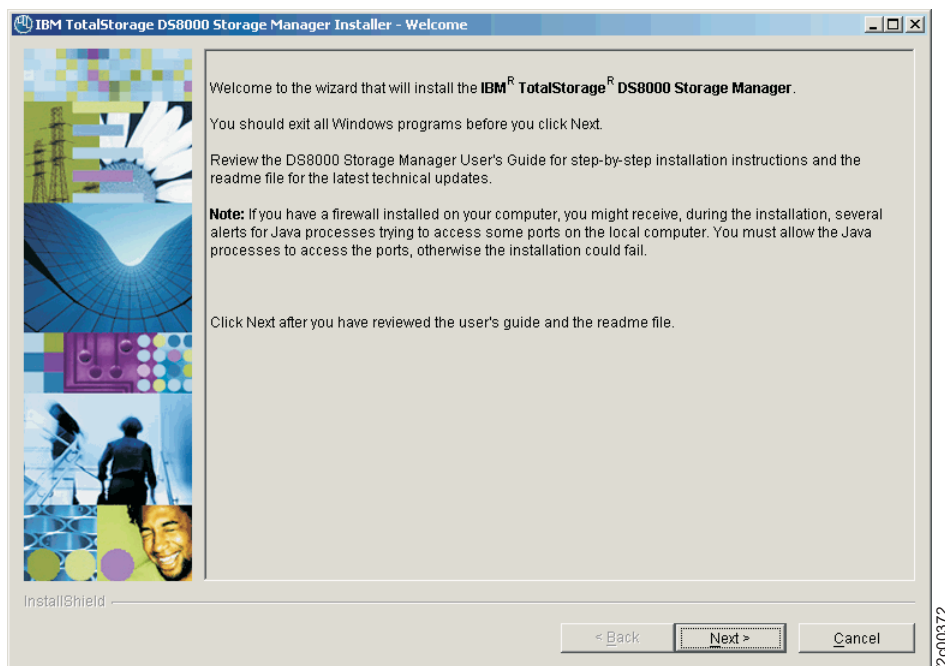
4. Click the **Readme file** selection on the LaunchPad to check for information that might supersede the information in this guide.
5. Click the **Installation wizard** selection on the LaunchPad to start the upgrade program.

Note: The LaunchPad window remains open behind the installation wizard so that you can access product information during the upgrade process.

There might be a slight delay while the software loads on your system. After the software loads, a DOS prompt window opens to display the following message:

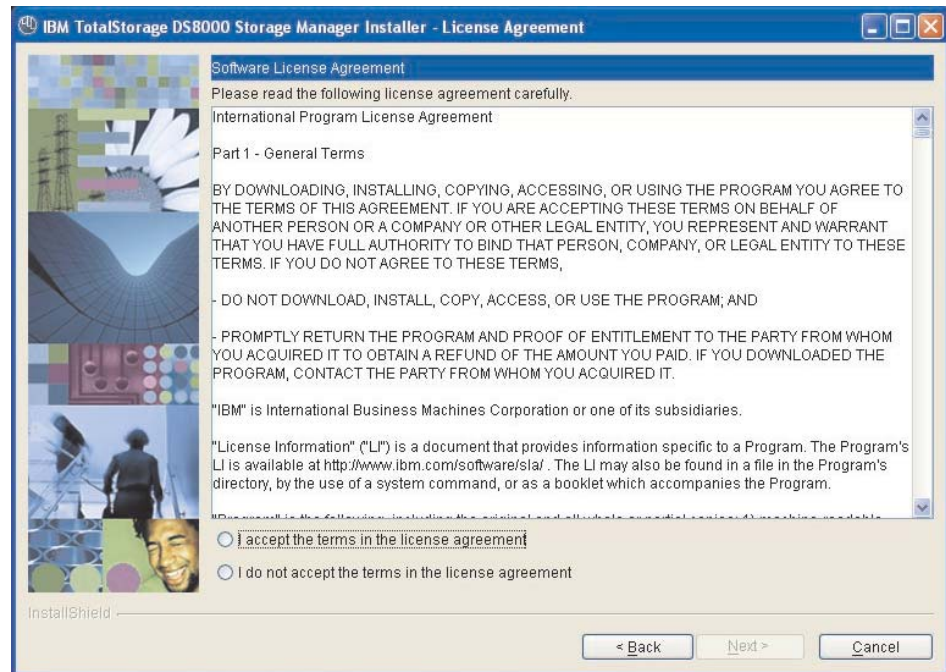
```
Initializing InstallShield Wizard...
Preparing Java (tm) Virtual Machine .....
.....
```

The Welcome window of the IBM TotalStorage DS Storage Manager upgrade program displays if no problems are discovered during the initial system check. If an error is discovered (for example, the operating system does not match the prerequisite), an error message is displayed and the upgrade program exits.



DS8000 Storage Manager Installer Welcome window

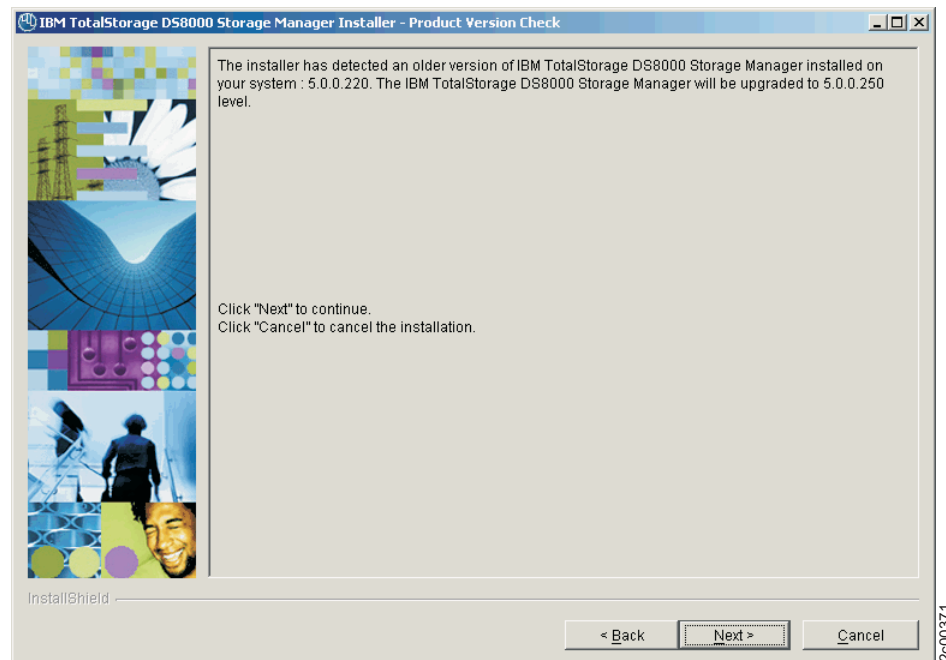
6. Click **Next** to continue, or click **Cancel** to exit the upgrade process. When you click Next, the License Agreement window displays.



License Agreement window

Note: Anytime you click **Cancel** on any of the upgrade windows, a message that asks you to confirm that you want to exit is displayed.

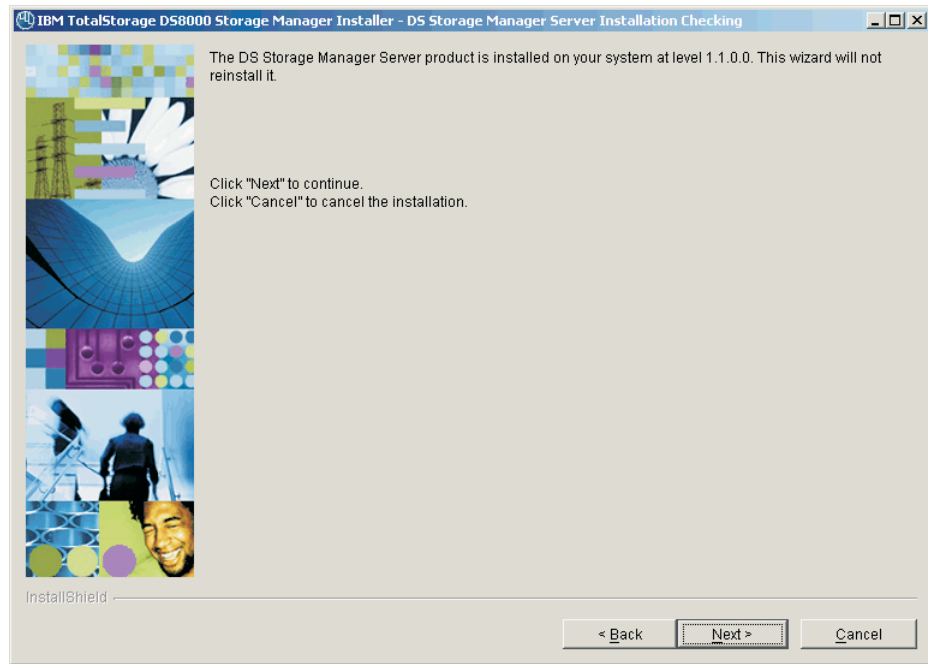
7. Read the license agreement and click your acceptance or nonacceptance of the agreement. If you accept, the **Next** button is highlighted. Click **Next** to continue or click **Cancel** to exit the upgrade process. When you click **Next**, the Product Version check window is displayed.



Product Version Check window

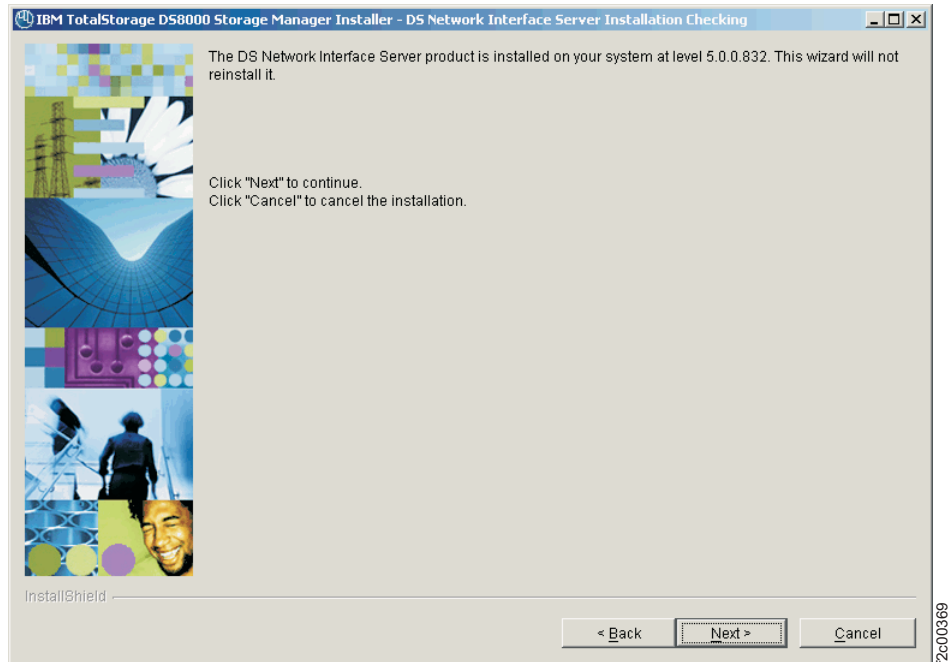
8. Observe the information that is displayed on the Product Version Check window regarding the version of the product that is already installed on your

system and the action that the wizard will perform. If you agree, click **Next** or click **Cancel** to exit the upgrade process. When you click **Next**, the DS Storage Manager Server Installation Checking window is displayed if the DS Storage Manager service is installed on your system.



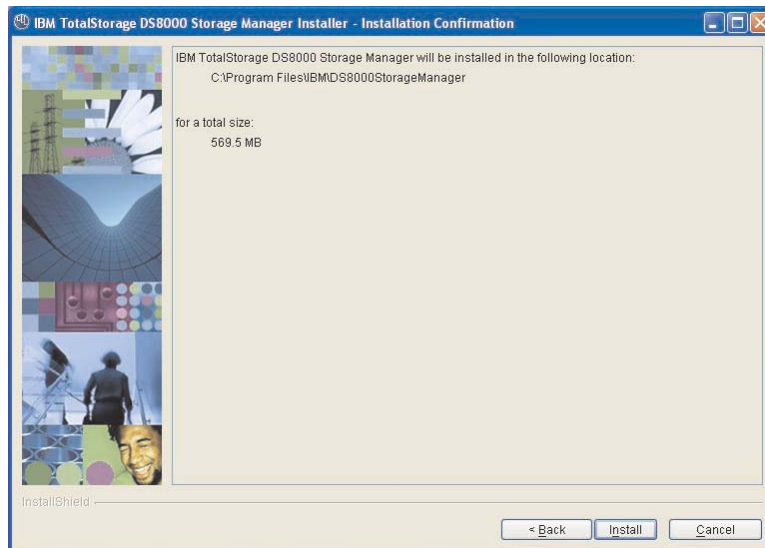
DS Storage Manager Server Installation Checking window

9. Observe the information that is displayed on the DS Storage Manager Server Installation Checking window regarding the version of the DS Storage Manager Server that is already installed on your system and the action that the wizard will perform. If you agree, click **Next** or click **Cancel** to exit the upgrade process. When you click **Next**, the DS Network Interface Server Installation Checking window is displayed if the DS Network Interface service is installed on your system.



DS Network Interface Server Installation Checking window

10. Observe the information that is displayed on the DS Network Interface Server Installation Checking window regarding the version of the DS Network Interface Server that is already installed on your system and the action that the wizard will perform. If you agree, click **Next** or click **Cancel** to exit the upgrade process. When you click **Next**, the Installation Confirmation window is displayed. This window displays both the location where the product will be installed and the total size that is required for the installation.



Windows Installation Confirmation window

11. Click the **Install** button on the Installation Confirmation window to begin the installation process. There are several progress windows that are displayed. There is no required interaction on your part for each of the progress windows that are displayed. However, you can choose to cancel (not recommended) the installation on any of the progress windows with varying consequences.

The installation process performs the following actions:

- a. If the two servers (DS Storage Manager Server and DS Network Interface Server) are already installed on your system, they are stopped in the following order of windows:
 - 1) The Embedded IBM WebSphere Application Server - Express server (part of the DS Storage Manager Server) window is stopped first.
 - 2) The service window (DS Network Interface Server) is stopped next.
 - 3) The WS Help System (part of the DS Storage Manager Server) window is stopped, if it was not stopped before by the Embedded IBM WebSphere Application Server.
- b. If one or both of the servers are not installed or have to be upgraded on your system, they are installed or upgraded in the following order (the progress of the installation is indicated on the associated progress window):
 - 1) DS Storage Manager Server Installation Progress window
 - 2) DS Network Interface Server Installation Progress window

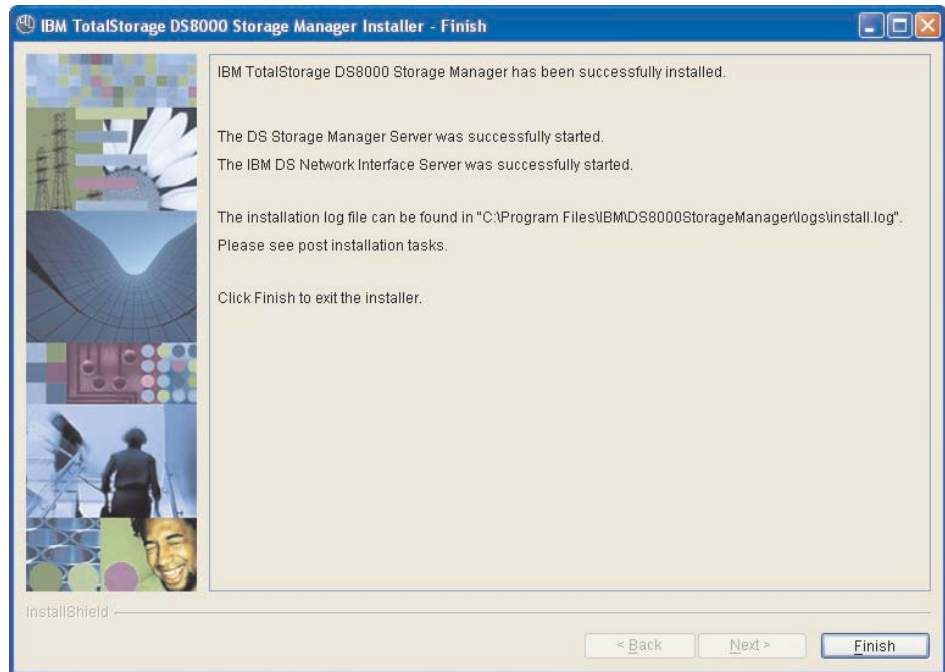
Note: You can click **Cancel** (not recommended) during the DS Network Interface Server installation process. The process does not stop immediately when the **Cancel** button is clicked. Rather, the process continues to install all the files that are associated with this part of the installation. These files remain installed and are not reinstalled upon reactivation of the installation process. A confirmation message is displayed asking you to affirm that you want to cancel the installation.

- c. The Components Installation Progress (displaying the installation or upgrade progress of the DS Storage Manager product applications) is displayed after the servers have been installed or upgraded.

Note: You can click **Cancel** (not recommended) during the components installation process. The installation process stops immediately when the **Cancel** button is clicked. A window with a confirmation message is displayed. For a new installation, when you confirm that you want to stop the process, all the files that have been copied up to the point that the **Cancel** button was clicked are uninstalled. You are then exited from the installation process.

When this part of the installation is completed, the system starts both servers: first the DS Network Interface Server and then the DS Storage Manager Server. Wait for the servers to be started before proceeding to the next step. When the servers have been started, the Installer Finish window is displayed.

- 12. Click **Finish** to exit the installation process. When the installation process exits, a file (postinstallation.txt) is displayed that contains the postinstallation tasks. Follow the instructions in this text file to complete the postinstallation tasks.



Windows DS8000 Storage Manager Installer Finish window

Note: If the installation fails, you must exit the installation process and check the install log for error messages.

13. If the Finish window indicates that the installation fails, check the installation log for error messages. The installation log is located in `xxx\logs\install.log`, where `xxx` is the destination directory where the IBM TotalStorage DS Storage Manager is installed (for example, `c:\Program Files\IBM\DS8000 Storage Manager`).

Errors fall into two categories: system or incorrect values that are designated during the installation process. Use the following guidelines to correct these errors.

- If the error is due to a system problem, correct it and reinstall the DS8000 Storage manager using either the interactive or silent mode of installation.
 - If the error is due to a wrong installation parameter value, restart the installation using the steps that are described in this procedure or the steps that are described in the silent mode installation process. Navigate to the Server Parameters window and insert the correct values. Finish the installation process.
14. Complete the postinstallation tasks. If, when exiting the installation wizard, the `postinstallation.txt` file does not automatically open, manually open it from the LaunchPad window, and follow the instructions to complete the postinstallation tasks.
 15. Exit the LaunchPad program by clicking **Exit** on the LaunchPad window.

Upgrading the IBM TotalStorage DS Storage Manager on the Windows operating system in unattended (silent) mode

Use the following steps to upgrade the IBM TotalStorage DS Storage Manager in your Windows environment using the unattended (silent) mode.

Note: The DS Storage Manager is not supported on any Windows 64-bit operating system.

The unattended (silent mode) upgrade option allows you to run the upgrade program unattended. Use this method of upgrade to issue a command from a command prompt window, in the root directory of the IBM TotalStorage DS Storage Manager CD.

Perform the following steps to upgrade the IBM TotalStorage DS Storage Manager in your Windows environment using the unattended mode:

1. Log on to your Windows system as an administrator.
2. Insert the IBM TotalStorage DS Storage Manager upgrade CD into the CD-ROM drive. If the interactive mode starts, click the **Exit** selection on the LaunchPad window to exit.
3. Type the following command at the command prompt and press the **Enter** key to start the upgrade process in silent mode: *setup.exe -silent*
4. Check the install log file for any possible error messages. This file is located in the **xxx\logs\install.log** directory, where xxx is the destination directory where the IBM TotalStorage DS Storage Manager is installed.
5. Start the IBM TotalStorage DS Storage Manager, if no errors are evident.
6. Perform the postinstallation tasks when the DS Storage Manager has been installed successfully. You can access the instructions for the postinstallation tasks from the postinstallation.txt file, in the doc directory on the product CD.

Chapter 3. Installing the IBM TotalStorage DS8000 Storage Manager on the Linux operating system

This section provides an overview of the installation, upgrade, and uninstall processes and instructions associated with the IBM TotalStorage DS8000 Storage Manager on the Linux operating system.

You can choose to install the IBM TotalStorage DS8000 Storage Manager on the Linux (Red Hat Enterprise Linux 2.1) operating system using either of the following modes:

- Graphical mode – allows you to use an online installation wizard that guides you through the installation process providing prompts and information needed to complete the installation.
- Unattended mode (also called silent mode) – allows you to customize a response file and issue a command to complete the installation process.

Installing the IBM TotalStorage DS Storage Manager on a Linux operating system using the graphical mode

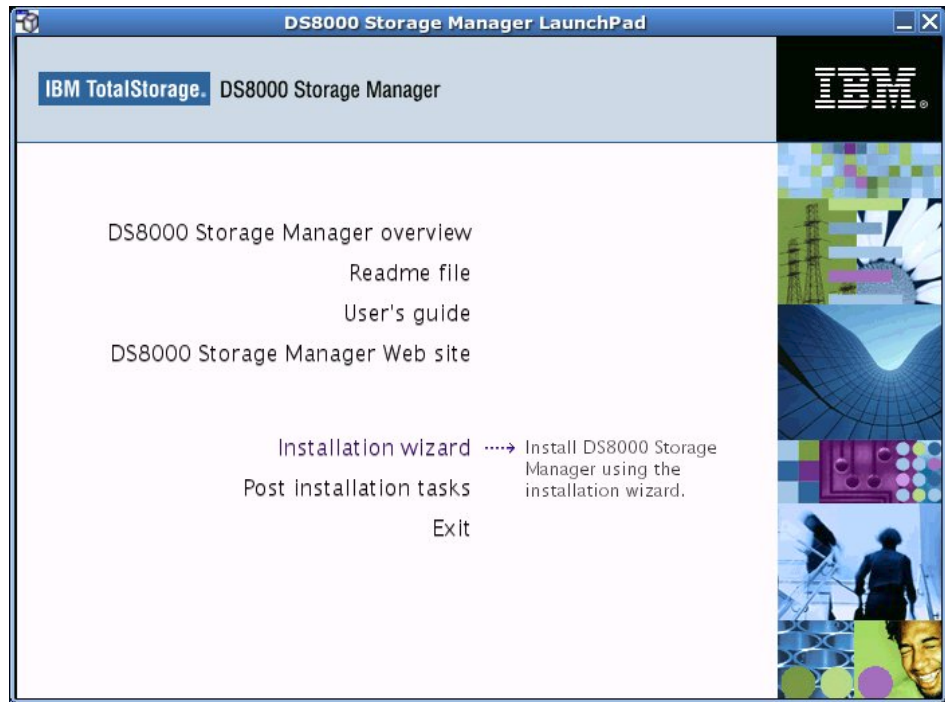
You can install the IBM TotalStorage DS Storage Manager in your Linux environment using the graphical mode.

Note: During the installation process, do not interrupt the program by abnormal methods, because an interruption can cause unpredictable effects in the system. Abnormal methods include using CTRL-C, closing the launchpad window which is open behind the installation windows, or shutting down the power source. To interrupt the installation process normally, use the Windows application **Cancel** button.

Perform the following steps to install the IBM TotalStorage DS Storage Manager.

1. Log on as a user with root authority.
2. Create a mount point or choose an existing mount point.
3. Determine where you want to mount the DS8000 Storage Manager CD. For example, imagine you have decided to mount the CD in a CD-ROM drive known as /dev/cdrom at mount point /mnt/cdrom. When the time comes for you to mount the CD in this drive, you perform the following steps:
 - a. Insert the CD in the CD-ROM drive.
 - b. Type the following command at the command prompt: **# mount -o ro /dev/cdrom /mnt/cdrom**
4. Choose the destination paths for the DS8000 Storage Manager components and be sure that you have enough space available on the chosen destination. The installation program (installer) provides a default installation location (**/opt/IBM/DS8000StorageManager**) for the component. You can override this by using the installer browse function. When you click **Next**, the installer analyzes the available space for the location that you provide. If there is not enough space, you cannot proceed to the next panel in the installation process.
5. Change the current directory to the mount point for the CD-ROM drive. For example, if the CD was mounted at the /mnt/cdrom mount point, type the following command: **# cd /mnt/cdrom**
6. Run the wizard launcher, launchpad_linux, by typing the following command: **# ./launchpad_linux**. The installation in graphical mode begins with a

LaunchPad facility to launch the installation program wizard.



Launchpad window

The LaunchPad facility provides links for you to view various text files, such as:

DS8000 Storage Manager overview

Provides information about the IBM TotalStorage DS Storage Manager software.

Readme file (recommended selection)

Provides last minute product information that was not provided in these installation instructions.

User's guide

Provides the instructions, in HTML format, associated with installing the DS8000 Storage Manager.

DS8000 Storage Manager Web site

Provides information from the product Web site.

Installation wizard

Starts the IBM TotalStorage DS Storage Manager installation program.

Post installation tasks

Provides information about configuring the IBM TotalStorage DS Storage Manager.

Exit

Exits the IBM TotalStorage DS Storage Manager LaunchPad program.

7. Click the **README** selection on the LaunchPad window to ensure that you have the latest installation information. The README file can provide additional information that supersedes information in this guide.

Note: You can also access the README.txt file located in the **doc** subdirectory on the IBM TotalStorage DS8000 Storage Manager CD.

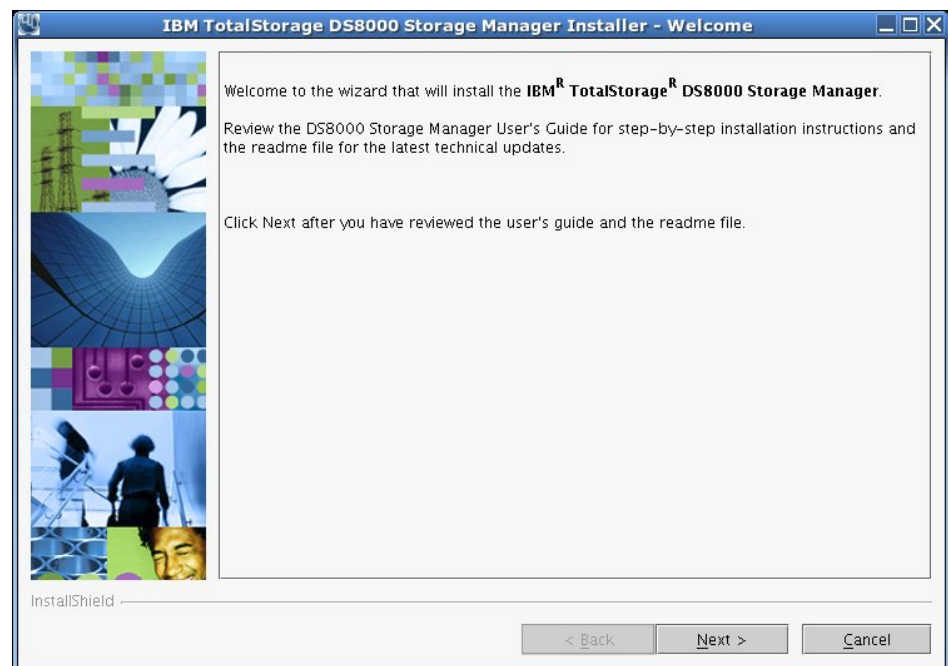
8. Click the **Installation wizard** selection on the LaunchPad to start the installation.

Note: The LaunchPad window remains open (behind the wizard) during the installation. You can access product information after the installation has started. The LaunchPad returns to the foreground when the installation is complete. You can click **Exit** to close the LaunchPad.

There might be a slight delay while the software loads on your system. After the software loads, a Command prompt window opens to display the following message:

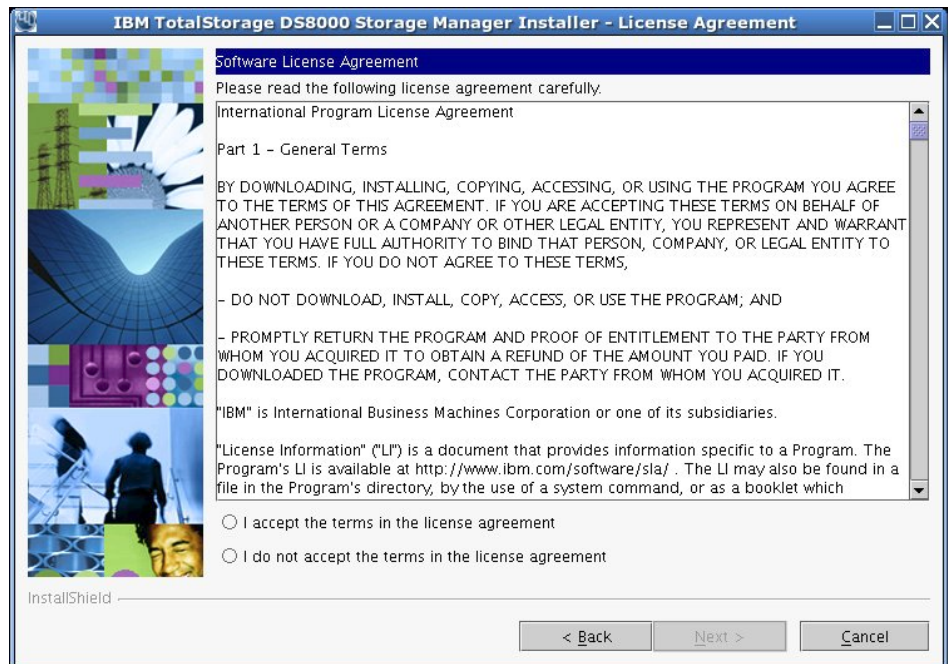
```
Initializing InstallShield Wizard...
Preparing Java (tm) Virtual Machine .....
.....
```

The Welcome window of the IBM TotalStorage DS Storage Manager installation program displays.



DS8000 Storage Manager Installer Welcome window

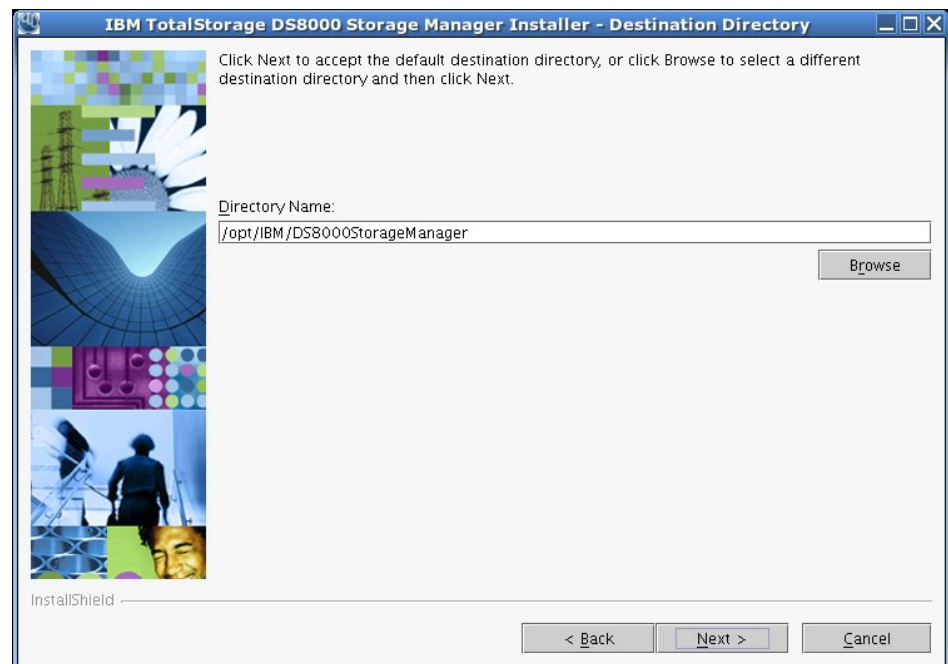
9. Click **Next** to continue, or click **Cancel** to exit the installation. When you click Next, the License Agreement window is displayed.



License Agreement window

Note: Anytime you click **Cancel** on any of the installation windows a message asking for you to confirm that you want to exit is displayed.

10. Read the license agreement and click your acceptance or nonacceptance of the agreement. If you accept, the **Next** button is highlighted. Click **Next** to continue or click **Cancel** to exit the installation. When you click **Next**, the Destination Directory window is displayed.



Destination Directory window

11. Choose the directory on the Destination Directory window where you want to install the application by using one of the following options. Clicking **Next** after any one of the options causes the server installation checking windows (a

window for each server) to display if the DS Storage Manager is installed on your system, or the Server Parameters window is displayed.

- a. Click the **Next** button to accept the default (recommended) directory that is shown in the Linux Destination Directory window.
- b. Type a fully qualified directory path into the Directory Name path to designate your own directory. Then click the **Next** button.
- c. Click the **Browse** button and a directory window is displayed. Scroll to the directory that you want to use for the application. Click it and click **OK**. The directory path is displayed in the Directory Name field. Click **Next** to continue.
- d. Click **Cancel** to exit the installation process.

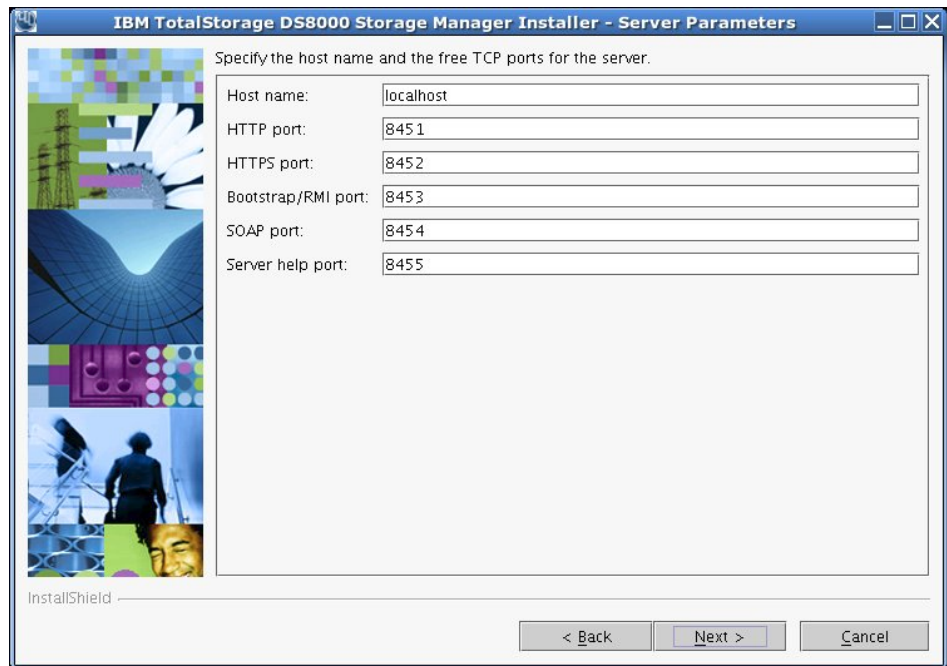
12. Observe the message on each of the next two server installation checking windows, and click **Next** on each window to continue. If the servers have not been installed, the Server Parameters window is displayed with default values in the input fields. Go to Step 13 to proceed with the Server Parameters window.

The DS Storage Manager Server Installation Checking window is displayed if the DS Storage Manager server is already installed on your system. A message is displayed on the window that indicates which version of the server is installed and whether the installation process installs a newer version of the server. To continue the installation process, you must click **Next**, or click **Cancel** to exit the installation process.

When you click **Next**, the DS Network Interface Server Installation Checking window is displayed with a message. This message indicates the version of the DS Network Interface server that is installed on your system and whether the installation process installs a newer version of the server. To continue the installation process, you must click **Next**, or click **Cancel** to exit the installation process.

When you click **Next**, the Installation Confirmation window is displayed. If your installation process has taken you through these two checking windows, go to Step 19 on page 61 to proceed with the installation.

13. Click **Next** after you ensure that all the fields on the Server Parameters window contain acceptable values. All fields must contain a value before you can proceed to the next step. You can change the default values by supplying the following information:



Linux Server Parameters window

- a. **Host Name** — The host name of the physical machine where the DS Storage Manager Server is installed. The host name must resolve to a physical network node on the server. When multiple network cards exist in the server, the host name or IP address must resolve to one of the network cards. The value that you specify for the host name is used as the value of the `hostName` property in WebSphere Application Server configuration files. Specify the host name value in one of the following formats:

- The fully qualified domain name servers (DNS) host name string; for example, *xmachine.manhattan.ibm.com*.
- The default short DNS host name string; for example, *xmachine*.
- The numeric IP address; for example, *9.127.255.3*.
- The generic name *localhost* or its equivalent IP address; for example, *127.0.0.1*.

Note: The specified value for the host name is stored by WebSphere Application Server into its internal configuration files. If you specified the host name in the fully qualified form or the short host name string, do not change the host name after the DS Storage Manager Server installation. If you do, the host name and the name that you specified for the configuration file no longer match. If you do change the host name, you must manually edit all the WebSphere Application Server configuration files that contain references to the old name to reflect the new name. If you specified the host name in the form of the numeric IP address, you must use a fixed IP address for the computer (you cannot use DHCP to obtain a dynamic IP address for it). The above restrictions do not exist if you specified the host name as *localhost*.

- b. **HTTP Port** — The port number that DS Storage Manager Server uses. The port number must be a port that is not being used by another process on the system. After DS Storage Manager Server is installed, you must include this port number in the Web address for opening the console. This

Web address is the protocol name, plus the fully qualified host name, plus the port, plus DS8000/Console. An example is
http://myhost.com:8451/DS8000/Console

- c. **HTTPS Port** — The port that DS Storage Manager Server uses for secure HTTP transport (HTTPS).
- d. **Bootstrap/RMI Port** — The port that DS Storage Manager Server uses for bootstrap port.
- e. **SOAP Port** — The port that DS Storage Manager Server uses for Simple Object Access Protocol (SOAP).
- f. **Server help port** — The port for the Eclipse help Server. This is the port that the help system (based on Eclipse technology) uses to receive requests for help files.

Note: The values used in the Server Parameters window fields must not conflict with existing port assignments that are already on the system or reserved for use by another application.

14. Click **Next** when the Server Parameters window is complete. The SSL Configuration window is displayed. This window allows you to supply the fully qualified name of the two server key that are files generated before or during the DS8000 Storage Manager installation.

IBM TotalStorage DS8000 Storage Manager Installer - SSL Configuration

Please enter the following information required to set SSL configuration for IBM TotalStorage DS8000 Storage Manager.

'Key file name' stands for the fully qualified path to the key file.
'Key file password' stands for the password for accessing the key file.
'Trust file name' stands for the fully qualified path to a trust file.
'Trust file password' stands for the password for accessing the trust file.

Select one of the following options concerning the SSL certificates:

☒ Generate the self-signed certificates during installation, or
☐ Use existing SSL certificates

Key file name:

Key file password:

Confirm key file password:

Trust file name:

Trust file password:

Confirm trust file password:

InstallShield

< Back Next > Cancel

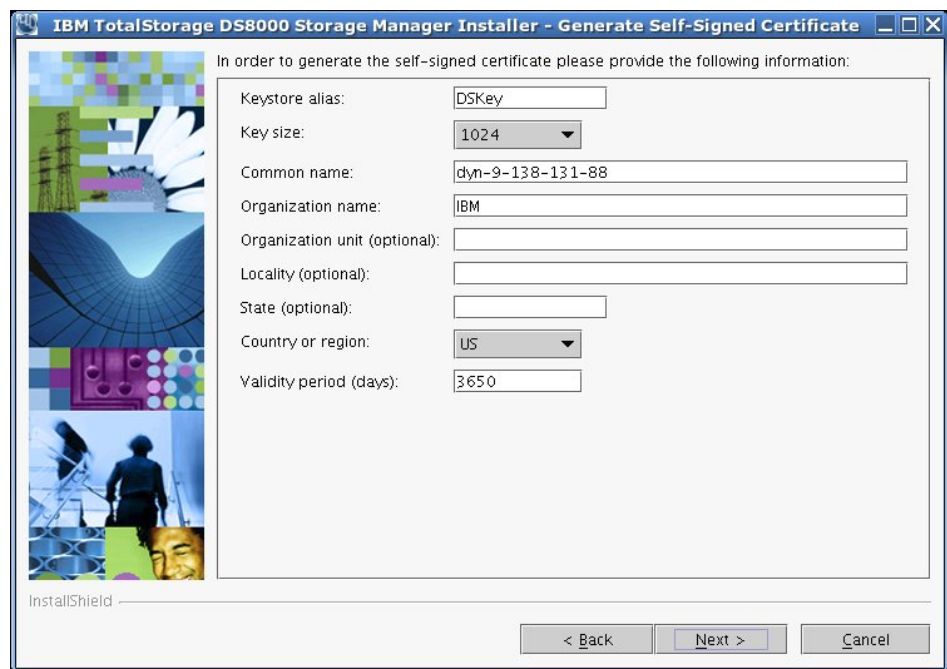
Linux SSL Configuration window

15. Follow the instructions on the window.
- a. First, designate whether you want to generate a new SSL certificate by selecting **Generate the self-signed certificates during installation**. This is the most likely selection when you are installing the DS8000 Storage Manager for the first time. You can also select **Use existing SSL certificates**. This choice is most likely made when the key files already exist and you do not want to generate new key files.
 - b. Complete the rest of the fields on the window. Provide the fully qualified path to the key file and trust file by using the Browse button. The key file

and trust file are identified by an extension of .jks. Also, provide a password for each file. The password must be a minimum of six characters and cannot contain any blanks.

16. Click **Next** to continue the installation process. If you have selected **Generate the self-signed certificates during installation**, the Generate Self-Signed Certificate window is displayed. If you selected **Use existing SSL certificates**, the Installation Confirmation window is displayed.

Note: When you select **Generate self-signed certificates during installation** and the key files already exist on the system at the specified location, clicking **Next** displays a message. The message asks you to designate whether you want to have the system delete existing application server certificates if they exist on the location and to generate others. The default is **NO**, which means that you want to keep the existing application server certificates.



IBM TotalStorage DS8000 Storage Manager Installer - Generate Self-Signed Certificate

In order to generate the self-signed certificate please provide the following information:

Keystore alias:	DSKey
Key size:	1024
Common name:	dyn-9-138-131-88
Organization name:	IBM
Organization unit (optional):	
Locality (optional):	
State (optional):	
Country or region:	US
Validity period (days):	3650

InstallShield

< Back Next > Cancel

Linux Generate Self-Signed Certificate window

17. Complete each of the input fields with the required information. Use the following as a guide to provide the needed information.

Keystore alias

This is the alias for the self-signed certificate. It uniquely identifies the certificate within the keystore file. It is a good practice to use a unique name that is related to the server name.

Key size

Two key sizes are available, 512 and 1024. 1024 is the default.

Common name

This name is the primary, universal identity for the certificate. If your system contains the secured WebSphere environment, this name must be valid in the configured user registry.

Organization name

The name of your company

Organization unit (optional)

The name of a department or division within your company. Remember that you are building a profile for the certificate. As a security measure the more specific the information the more secure your system.

Locality (optional)

The city or location where your company resides.

State (optional)

The state or province where your company resides.

Country or region

The two-character designator that identifies the country where your company is located. Use the pulldown menu to make this selection.

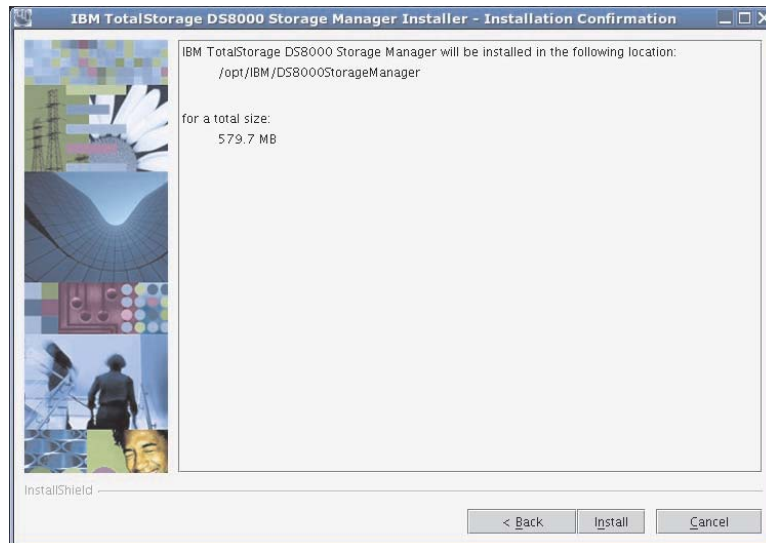
Validity period (days)

Specifies the lifetime of the certificate in days.

The information that you provide is used to build a profile for the certificate. This information is used during identity processing to ensure validation of any processing that is occurring. It is part of the security measures that is used during SSL connectivity.

Click **Next** to continue, or click **Cancel** to exit the installation.

18. Click **Next**. The Installation Confirmation window is displayed. This window displays both the location where the product will be installed and the total size that is needed for the installation

**Linux Installation Confirmation window**

19. Click the **Install** button on the Installation Confirmation window to begin the installation process. There are several progress windows that are displayed. There is no required interaction on your part for each of the progress windows that are displayed. However, you can choose to cancel (not recommended) the installation on any of the progress windows with varying consequences.

The installation process performs the following actions:

- a. If the two servers (DS Storage Manager Server and DS Network Interface Server) are already installed on your system, they are stopped in the following order of windows:
 - 1) The Embedded IBM WebSphere Application Server - Express server (part of the DS Storage Manager Server) window is stopped first.
 - 2) The service window (DS Network Interface Server) is stopped next.

- 3) The WS Help System (part of the DS Storage Manager Server) window is stopped, if it was not stopped before by the Embedded IBM WebSphere Application Server.
- b. If one or both of the servers are not installed or have to be upgraded on your system, they are installed or upgraded in the following order (the progress of the installation is indicated on the associated progress window):
 - 1) DS Storage Manager Server Installation Progress window
 - 2) DS Network Interface Server Installation Progress window

Note: You can click **Cancel** (not recommended) during the DS Network Interface Server installation process. The process does not stop immediately when the **Cancel** button is clicked. Rather, the process continues to install all the files that are associated with this part of the installation. These files remain installed and are not reinstalled upon reactivation of the installation process.

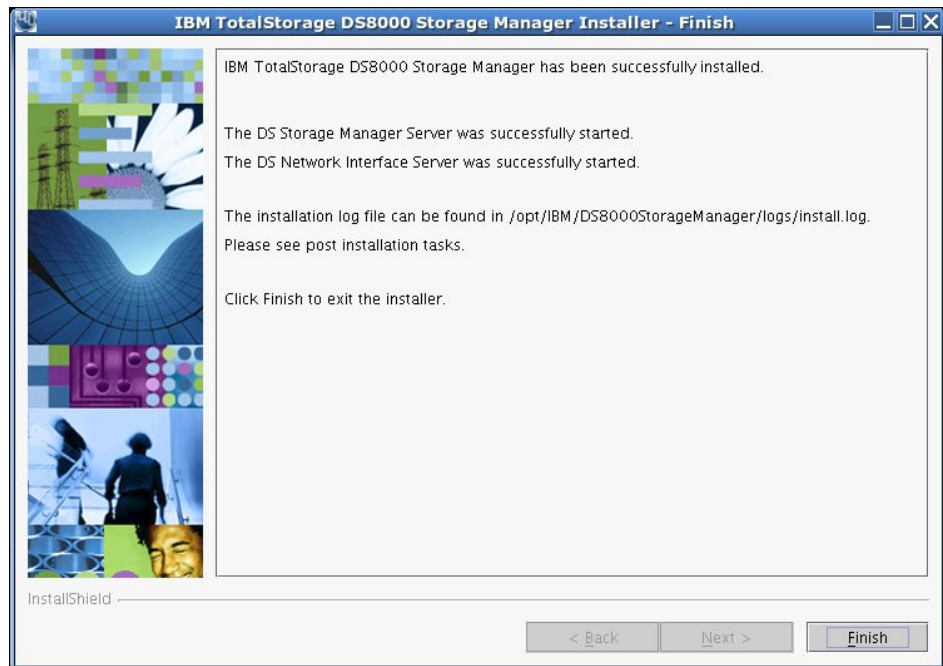
A confirmation message is displayed that asks you to confirm that you want to cancel the installation.

- c. The Components Installation Progress (displaying the installation or upgrade progress of the DS Storage Manager product applications) is displayed after the servers have been installed or upgraded.

Note: You can click **Cancel** (not recommended) during the components installation process. The installation process stops immediately when the **Cancel** button is clicked. A window with a confirmation message is displayed. For a new installation, when you confirm that you want to leave, all the files that have been copied up to the point that the **Cancel** button was clicked are uninstalled. You then are exited from the installation process.

When this part of the installation is completed, the system starts both servers: first the DS Network Interface Server and then the DS Storage Manager Server. Wait for the servers to be started before going to the next step. When the servers have been started, the Installer Finish window is displayed.

- 20. Click **Finish** to exit the installation process. When the installation wizard exits, a file (postinstallation.txt) is displayed, containing the postinstallation tasks. Follow the instructions in the postinstallation text file to complete the postinstallation tasks. If the installation failed, you must exit the installation process and check the install log for error messages.



Linux DS8000 Storage Manager Installer Finish window

21. If the Finish window indicates that the installation failed, check the installation log for error messages. The installation log is located in *xxx/logs/install.log*, where *xxx* is the destination directory where the IBM TotalStorage DS Storage Manager is installed.
Errors fall into two categories: system or incorrect values that are designated during the install process. Use the following guidelines to correct these errors:
 - If the error is due to a system problem, correct it and reinstall the GUI using either the graphical or silent mode of install.
 - If the error is due to a wrong installation parameter value, restart the installation using the steps that are described in this procedure or the steps that are described in the silent mode install. Navigate to the Console Server Parameters window and insert the correct values. Then finish the installation.
22. Check the postinstallation instructions after the installation has completed. The *postinstallation.txt* file is displayed immediately after the installation of the DS Storage Manager finishes processing. If the *postallation.txt* file does not display, click the **Post installation tasks** link on the LaunchPad window to access these instructions. These instructions provide reminders and the steps to help you start using the DS Storage Manager.
23. Exit the LaunchPad program by clicking **Exit** on the LaunchPad window.
24. Type the following command when you are finished using the IBM TotalStorage DS Storage Manager CD and want to remove it from your computer: `# umount /mnt/cdrom` where */mnt/cdrom* is where the CD is mounted.

Installing the IBM TotalStorage DS8000 Storage Manager on the Linux operating system in unattended (silent) mode

You can install the IBM TotalStorage DS Storage Manager in your Linux environment using the unattended (silent) mode.

The unattended (silent) installation option allows you to run the installation program unattended. Use this method of installation to customize a response file and to issue a command from a command prompt window. The response file is a template on the IBM TotalStorage DS Storage Manager CD.

Perform the following steps to install the IBM TotalStorage DS Storage Manager in your Linux environment using the unattended mode:

1. Log on as a user with root authority.
2. Create a mount point or choose an existing mount point.
3. Determine where you want to mount the DS8000 Storage Manager CD. For example, imagine you have decided to mount the CD in a CD-ROM drive known as `/dev/cdrom` at mount point `/mnt/cdrom`. When the time comes for you to mount the CD in this drive, perform the following steps:
 - a. Insert the CD in the CD-ROM drive.
 - b. Type the following command at the command prompt: **# mount -o ro /dev/cdrom /mnt/cdrom**
4. Choose the destination paths for the DS8000 Storage Manager components and be sure you have enough available space on the chosen destination. The installation program (installer) provides a default installation location (**/opt/IBM/DS8000StorageManager**) for the component. You can override this by using the `product.installLocation` parameter (option) from the response file used by the silent installation.
5. Change the current directory to the mount point for the CD-ROM drive. For example, if the CD was mounted at the `/mnt/cdrom` mount point, type the following command: **# cd /mnt/cdrom**
6. Read the `README.txt` file located on the DS8000 Storage Manager CD to check for information that might supersede the information in these installation instructions.
7. Copy the template response file (`responsefile.txt`) from the root directory of the product CD to a location on your system in order to modify it according to your installation parameters. The response file contains all the installation parameters whose values can be modified during an interactive installation.
8. Modify your copy of the response file in accordance with the following guidelines:
 - a. Open the response file using a text editor. Remember that you *must* enclose all values in double quotation marks ("`<value>`").
 - b. Replace the parameters values with the appropriate values for each uncommented option.
 - c. For the port values, ensure that they are free for use (not used or reserved for use by other applications on your system). Use the command **netstat -an** to verify the ports that are in use on your system.
 - d. You must specify two passwords for the SSL keys files. These passwords must observe the following criteria:
 - The passwords must contain at least 6 characters.
 - Use the following options within the `responsefile.txt` file to set these options:
 - For the keystore use:

`-W wasSslConfiguration.keystorePassword="your_keystore_password"`
 - For the truststore use:

```
-W wasSslConfiguration.truststorePassword="<your_truststore_password>"
```

Note: These options do not have default values. If you do not set the values, the installation process fails.

e. Save the changes made to the response/options file.

9. Type **`./setup.linux -options <responsefile-path>/<responsefile>`** in the same Command Prompt window in which you changed the current directory to the root of the CD, to start the DS8000 Storage Manager installation program in silent mode.

<responsefile-path>

Designates the path where the response file was created before

<responsefile>

Designates the relative name of the response file that you created and modified before

10. Wait for the DS8000 Storage Manager installation program to install the product. This can take 5 - 10 minutes.

Note: During this processing, the installation process checks to see if the DS Storage Manager server and DS Network Interface Server are already installed and if they are the current version or need to be upgraded. If they are already installed and need to be upgraded, the installation process performs this function using the already set values from the previous installation, found on the system in the server configuration files, for all the associated values. The specified values in the response file are ignored.

11. Review the log file for any possible error messages, after the installation program completes. If there are no errors, the installation program completed successfully. The log file is located in `xxx/logs/install.log`, where `xxx` is the destination directory where the DS8000 Storage Manager was installed. The `install.log` file contains a trace of the installation actions.
12. Perform the postinstallation tasks after the installation has completed. You can read the instructions for the postinstallation tasks in the `postinstallation.txt` in the `doc` directory on the product CD.

The following is an example of the template response file that is created when you process the previous steps.

```
#####
#
# InstallShield Options File
#
# Wizard name: Setup
# Wizard source: setup.jar
# Created on: Tue Jul 27 09:02:47 EEST 2004
# Created by: InstallShield Options File Generator
#
# This file contains values that were specified during a recent execution of
# Setup. It can be used to configure Setup with the options specified below when
# the wizard is run with the "-options" command line option. Read each setting's
# documentation for information on how to change its value.
#
# A common use of an options file is to run the wizard in silent mode. This lets
# the options file author specify wizard settings without having to run the
# wizard in graphical or console mode. To use this options file for silent mode
# execution, use the following command line arguments when running the wizard:
#
# -options "responsefile.txt"
#
#####

-silent

#####
#
# IBM TotalStorage DS8000 Storage Manager Install Location
#
# The install location of the product. Specify a valid directory into which the
# product should be installed. If the directory contains spaces, enclose it in
# double-quotes.
#

-P product.installLocation="/opt/IBM/DS8000StorageManager"

#####
#
# User Input Field - hostname
#
# The fully-qualified host name of the machine where DS Storage Manager Server
# will be installed.
#

-W wasExpressConfig.hostname="localhost"
#####
#
# User Input Field - httpPort
#
# The port number that the HTTP transport in the application server will use.
# This value must not conflict with existing port
# assignments on the system.
#

-W wasExpressConfig.httpPort="8451"

#####
#
# User Input Field - httpsPort
#
# The port number that the HTTPS transport in the application server will use
# for secure HTTP transport. This value must not conflict with existing port
# assignments on the system.

-W wasExpressConfig.httpsPort="8452"

#####
#
```

```

# User Input Field - bootstrapPort
#
# The address for the bootstrap function and the port number for the Java Remote
# Method Invocation (RMI) connector in the application server. This value must
# not conflict with existing port assignments on the system.
#

-W wasExpressConfig.bootstrapPort="8453"

#####
#
# User Input Field - soapPort
#
# The address for the Simple Object Access Protocol (SOAP) connector in the
# application server. This value must not conflict with existing port
# assignments on the system.
#

-W wasExpressConfig.soapPort="8454"

#####
#
# User Input Field - helpPort
#
# The port for Eclipse
# This value must not conflict with existing port assignments on the system.
#

-W wasExpressConfig.helpPort="8455"

#####
#
# User Input Field - certificateOption
#
# Select one of the following options concerning the SSL
# certificates. Legal values are: "generate" (the default value) - Generate the
# self-signed certificates during installation. "use" - Use existing
# certificatesFor example, to specify that the "generate" option is selected,
# use -W wasSslConfigurationPanel.certificateOption="generate"
#

-W wasSslConfiguration.certificateOption="generate"

#####
#
# User Input Field - keystoreFileName
#
# The absolute path of the keystore file.
#

-W wasSslConfiguration.keystoreFileName="/opt/IBM/
SMSEServer/keys/SMSEServerKeyFile.jks"

#####
#
# User Input Field - keystorePassword
#
# The password for the keystore file.
#

-W wasSslConfiguration.keystorePassword=""

#####

```

```

#
# User Input Field - truststoreFileName
#
# The absolute path of the truststore file.
#

-W wasSslConfiguration.truststoreFileName="/opt/IBM/
SMServer/keys/SMServerTrustFile.jks"

#####
#
# User Input Field - truststorePassword
#
# The password for the truststore file.
#

-W wasSslConfiguration.truststorePassword="<your_truststore_password>"

#####
#
# Delete server certificates option
#
# Option to delete certificates files if they exist. Legal values are: "yes" - in
# order to delete application server certificates if they exist on the location,
# and to generate others, or "no" (the default value) - in order to keep the
# existing application server certificates.
#

-G deleteCertificates=yes

#####
#
# User Input Field - keystoreAlias
#
# The alias for the self-signed digital certificate, which is used to uniquely
# identify the certificate within the keystore file. If you have only one
# certificate in each keystore file, you can assign any value to the label.
# However, it is good practice to use a unique label related to the server name.
#

-W CertificateParams.keystoreAlias="DSKey"

#####
#
# User Input Field - keySize
#
# The key size. That must be between 512 and 1024 and must be multiple of
# 64b. The only two allowed values are 512 or 1024. If you enter other values it
# takes the default value 1024.
#

-W CertificateParams.keySize="1024"

#####
#
# User Input Field - commonName
#
# The common name is the primary, universal identity for the certificate; it
# should uniquely identify the principal that it represents. In a WebSphere
# environment, certificates frequently represent server principals, and the
# common convention is to use common names of the form "host_name" and
# "server_name". The common name must be valid in the configured user registry
# for the secured WebSphere environment.
#

```

```

-W CertificateParams.commonName=""

#####
#
# User Input Field - organizationName
#
# The name of your organization.
#

-W CertificateParams.organizationName="IBM"

#####
#
# User Input Field - organizationUnit
#
# The organization unit (a department or division). For a self-signed
# certificate, these fields are optional. However, commercial CAs might require
# them.
#

-W CertificateParams.organizationUnit=""

#####
#
# User Input Field - localityName
#
# The location (city). For a self-signed certificate, these fields are optional.
# However, commercial CAs might require them.
#

-W CertificateParams.localityName=""

#####
#
# User Input Field - state
#
# The state or province (if applicable). For a self-signed certificate, these
# fields are optional. However, commercial CAs might require them.
#

-W CertificateParams.state=""

#####
#
# User Input Field - country
#
# The two-letter identifier of the country/region in which the server belongs.
# For a self-signed certificate, these fields are optional. However, commercial
# CAs might require them.
#

-W CertificateParams.country="US"

#####
#
# User Input Field - validity
#
# The lifetime of the certificate in days.
#

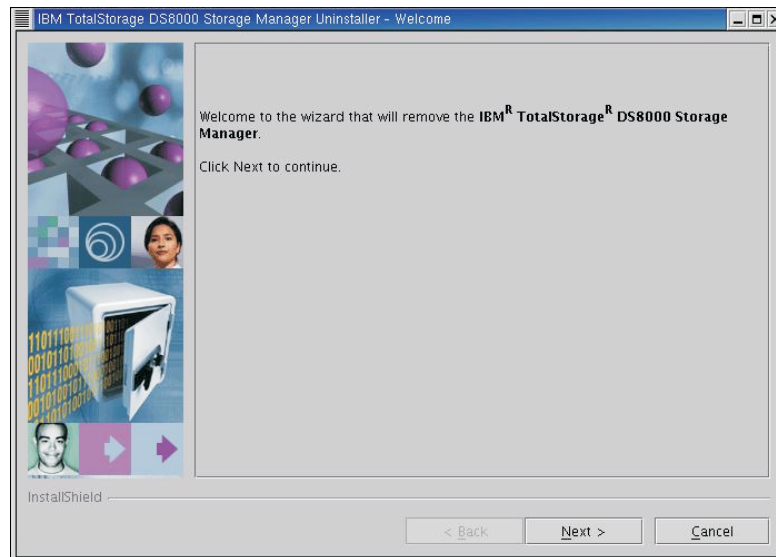
-W CertificateParams.validity="3650"

```

Removing the IBM TotalStorage DS Storage Manager from a Linux operating system

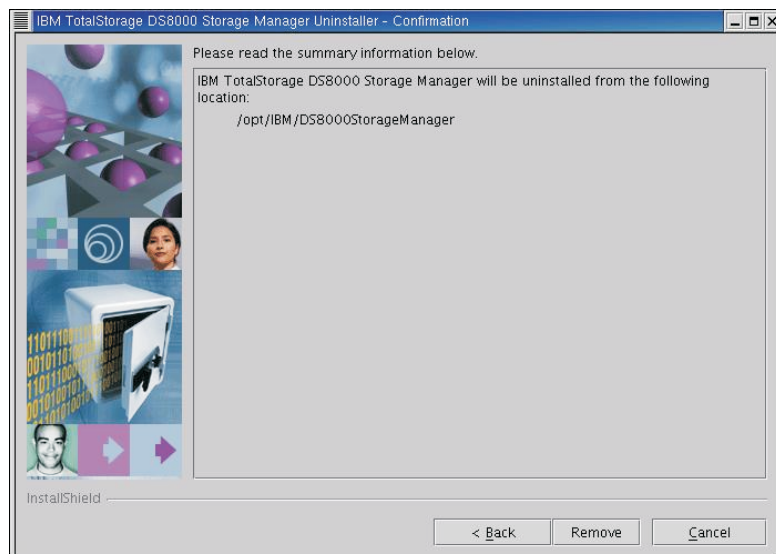
You can remove the IBM TotalStorage DS Storage Manager by following the steps in this procedure.

1. Log on as root to the system where the DS8000 Storage Manager is installed.
2. Type: # `<dest_path>/_uninst/uninstall` to run the uninstallation program, `uninstall`, from the `_uninst` subdirectory of `<dest_path>`. `<dest_path>` is the path where the IBM TotalStorage DS Storage Manager is installed. If you used the default path, it is `/opt/IBM/DS8000StorageManager`. The Uninstaller Welcome window is displayed.



Uninstaller Welcome window

3. Click **Next** to continue or click **Cancel** to exit the removal process. When you click **Next**, the Confirmation window is displayed showing the directory from which the IBM TotalStorage DS Storage Manager GUI program will be removed.



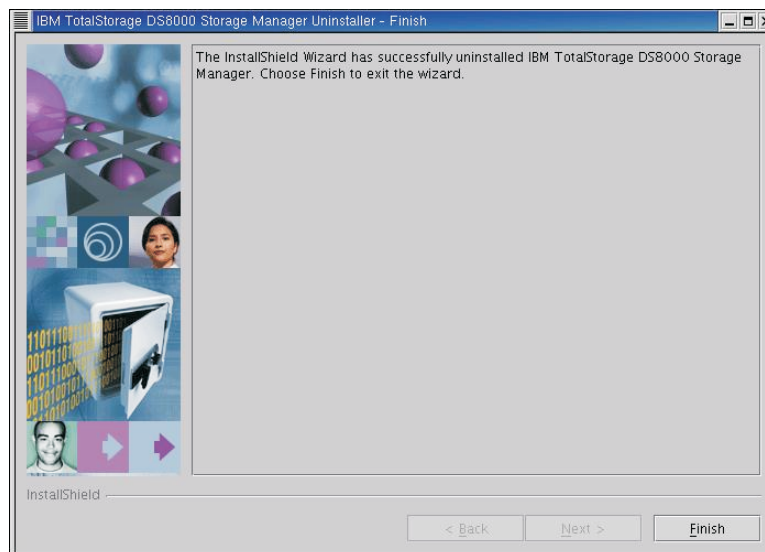
Uninstaller Linux Confirmation window

4. Click **Remove** to continue or **Cancel** to stop the removal and exit the uninstallation process. Click **Back** to return to the previous window. When you click **Remove**, the Uninstaller Uninstallation Progress window is displayed. When the uninstallation process is finished, the Finish window is displayed containing a statement about the success or failure of the uninstallation process.

During the uninstallation process, the following system checks are made:

- A check of your system to detect if DS Storage Manager Server is used by other programs on the system and if not used, removes it from the system. Otherwise, the DS Storage Manager Server is not removed from your system.
- A check of your system to detect if DS Network Interface Server is used by other programs on the system and if not used, removes it from the system. Otherwise, the DS Network Interface Server is not removed from your system.

When the uninstall process is finished, the Uninstaller Finish window is displayed containing a statement about the success or failure of the uninstallation process.



Uninstaller Finish window

5. Click **Finish** to complete the removal process and exit the uninstallation program.

Removing the DS8000 Storage Manager from a Linux operating system using the silent mode

You can remove the IBM TotalStorage DS Storage Manager from your Linux operating system by using the silent (unattended mode).

1. Log on as root to the system where the DS8000 Storage Manager is installed.
2. Type `# <dest_path>_uninst\uninstall.exe -silent`, where `<dest_path>` is the path where the DS8000 Storage Manager is installed. This runs the uninstallation program.
3. Wait for the program to remove the DS8000 Storage Manager.

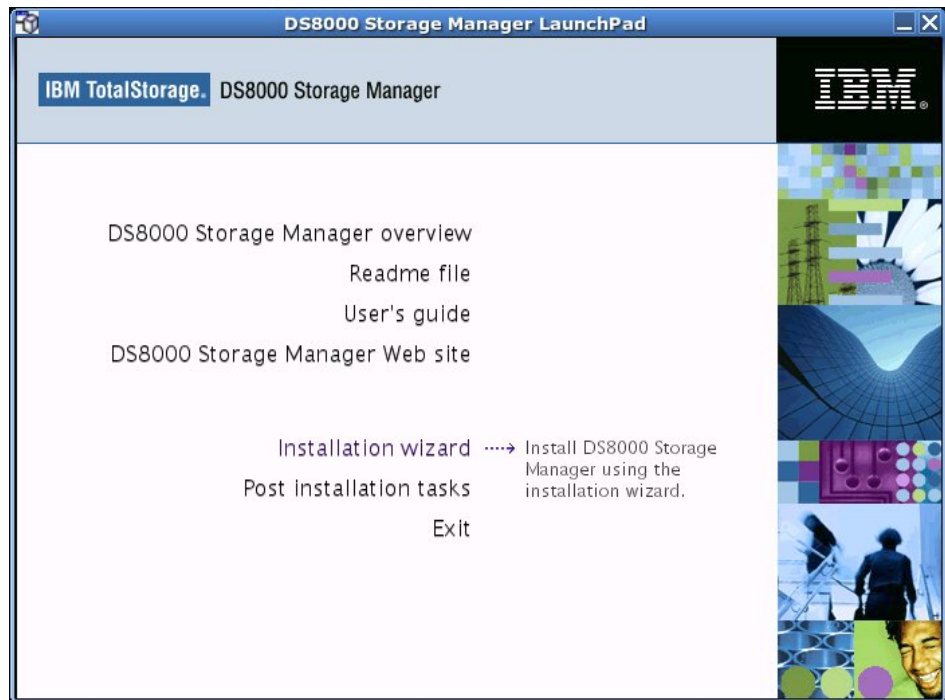
Upgrading the IBM TotalStorage DS Storage Manager on a Linux operating system

Use this process to upgrade the DS Storage Manager to the latest level.

Note: During the upgrade process, do not interrupt the program by abnormal methods, because an interruption can cause unpredictable effects in the system. Abnormal methods include using CTRL-C, closing the launchpad window which is open behind the upgrade windows, or shutting down the power source. To interrupt the upgrade process normally, use the **Cancel** button displayed on an upgrade window.

Perform the following steps to upgrade the IBM TotalStorage DS Storage Manager.

1. Log on as a user with root authority.
2. Create a mount point or choose an existing mount point.
3. Determine where you want to mount the DS8000 Storage Manager CD. For example, imagine you have decided to mount the CD in a CD-ROM drive known as /dev/cdrom at mount point /mnt/cdrom. When the time comes for you to mount the CD in this drive, perform the following steps:
 - a. Insert the CD in the CD-ROM drive.
 - b. Type the following command at the command prompt: **# mount -o ro /dev/cdrom /mnt/cdrom**
4. Choose the destination paths for the DS8000 Storage Manager components and be sure that you have enough space available on the chosen destination. The upgrade program (installer) provides a default installation location (**/opt/IBM/DS8000StorageManager**) for the component. You can override the default installation location by using the installer browse function. When you click **Next**, the installer analyzes the available space for the location that you provide. If there is not enough space, you cannot proceed to the next panel in the upgrade process.
5. Change the current directory to the mount point for the CD-ROM drive. For example, if the CD was mounted at the /mnt/cdrom mount point, type the following command: **# cd /mnt/cdrom**.
6. Run the wizard launcher, launchpad_linux, by typing the following command: **# ./launchpad_linux**. The upgrade in graphical mode begins with a LaunchPad facility to launch the upgrade program wizard.



Launchpad window

The LaunchPad facility provides links for you to view various text files, such as:

DS8000 Storage Manager overview

Provides information about the IBM TotalStorage DS Storage Manager software.

Readme file (recommended selection)

Provides last minute product information that was not provided in these upgrade instructions.

User's guide

Provides the instructions, in HTML format, associated with installing and upgrading the DS8000 Storage Manager.

DS8000 Storage Manager Web site

Provides information from the product Web site.

Installation wizard

Starts the IBM TotalStorage DS Storage Manager upgrade program.

Post installation tasks

Provides information about configuring the IBM TotalStorage DS Storage Manager.

Exit

Exits the IBM TotalStorage DS Storage Manager LaunchPad program.

7. Click the **README** selection on the LaunchPad window to ensure that you have the latest installation and upgrade information. The README file can provide additional information that supersedes information in this guide.

Note: You can also access the README.txt file located in the **doc** subdirectory on the IBM TotalStorage DS8000 Storage Manager CD.

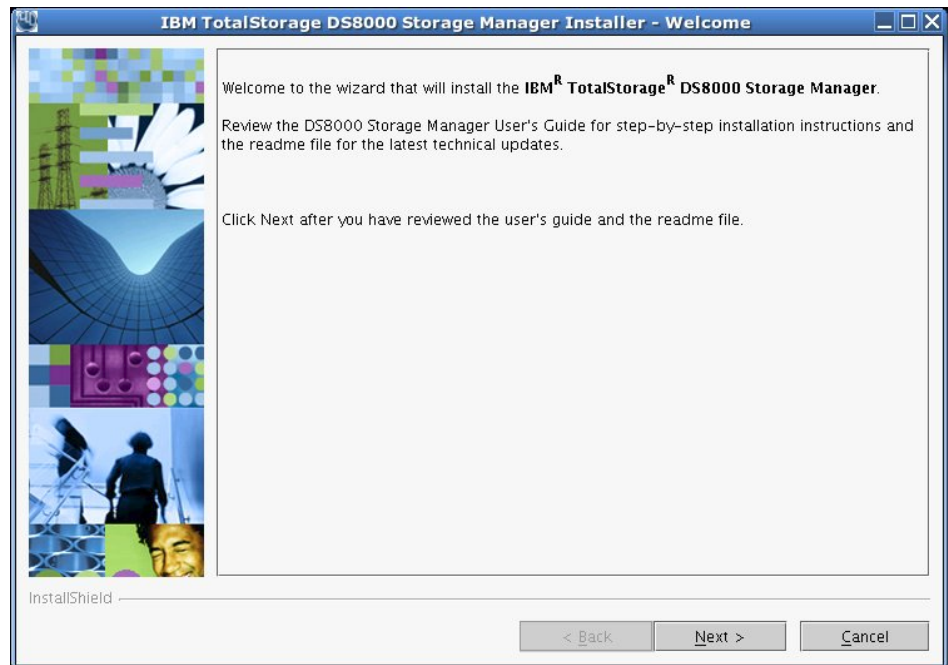
8. Click the **Installation wizard** selection on the LaunchPad to start the upgrade.

Note: The LaunchPad window remains open (behind the wizard) during the upgrade. You can access product information after the upgrade has started. The LaunchPad returns to the foreground when the upgrade is complete. You can click **Exit** to close the LaunchPad.

There might be a slight delay while the software loads on your system. After the software loads, a Command prompt window opens to display the following message:

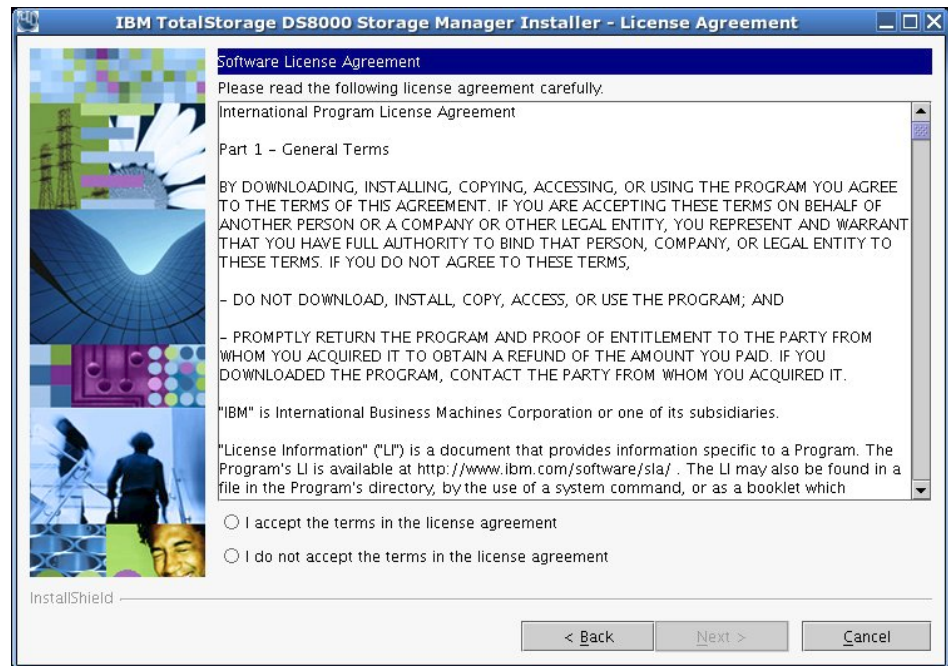
```
Initializing InstallShield Wizard...
Preparing Java (tm) Virtual Machine .....
.....
```

The Welcome window of the IBM TotalStorage DS Storage Manager installation program displays.



DS8000 Storage Manager Installer Welcome window

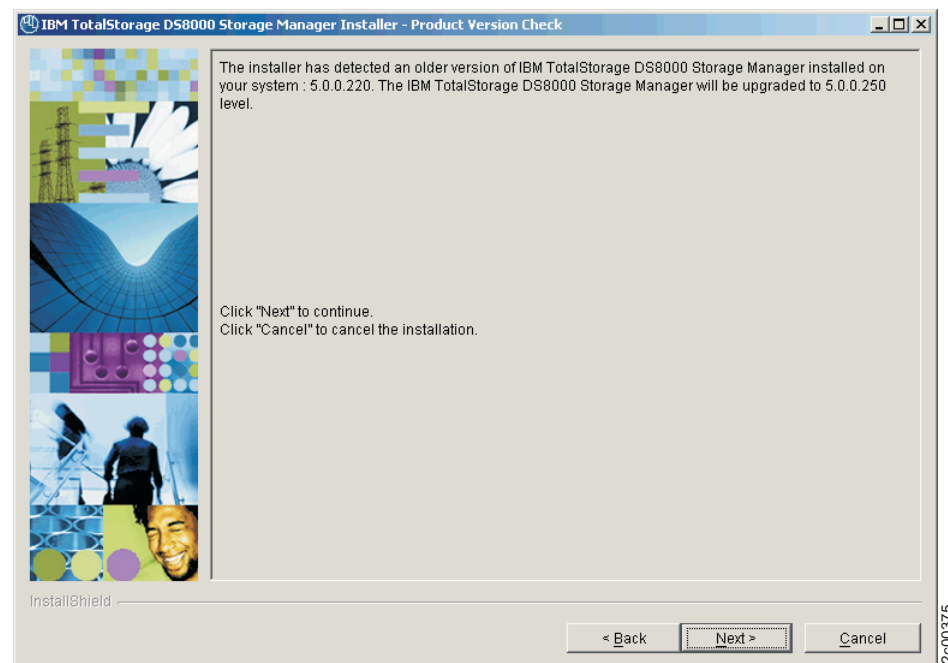
9. Click **Next** to continue, or click **Cancel** to exit the upgrade process. When you click **Next**, the License Agreement window is displayed.



License Agreement window

Note: Anytime you click **Cancel** on any of the upgrade windows a message that asks you to confirm that you want to exit is displayed.

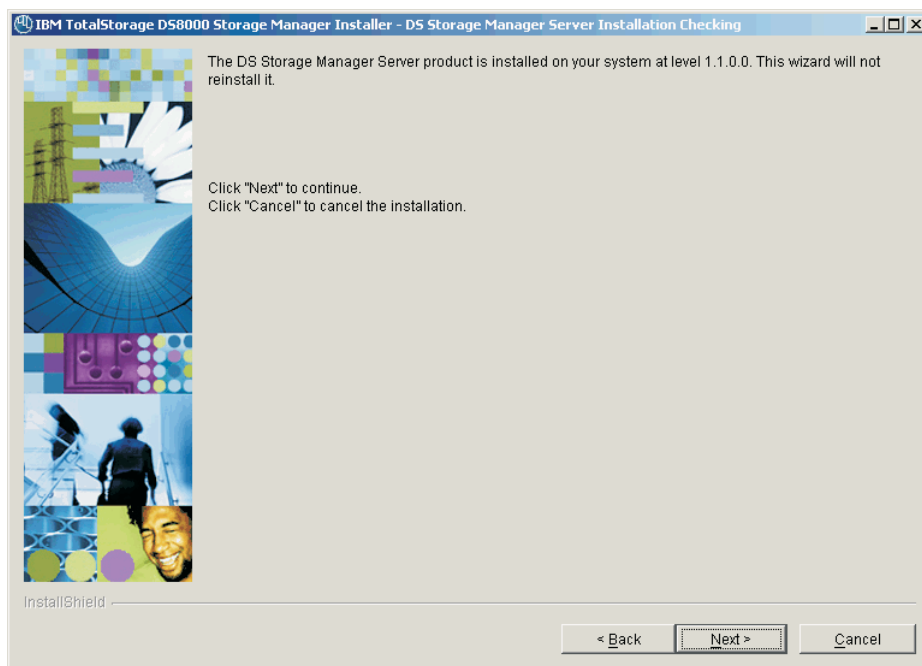
10. Read the license agreement and click your acceptance or nonacceptance of the agreement. If you accept, the **Next** button is highlighted. Click **Next** to continue or click **Cancel** to exit the upgrade process. When you click **Next**, the Product Version Check window is displayed.



Product Version Check window

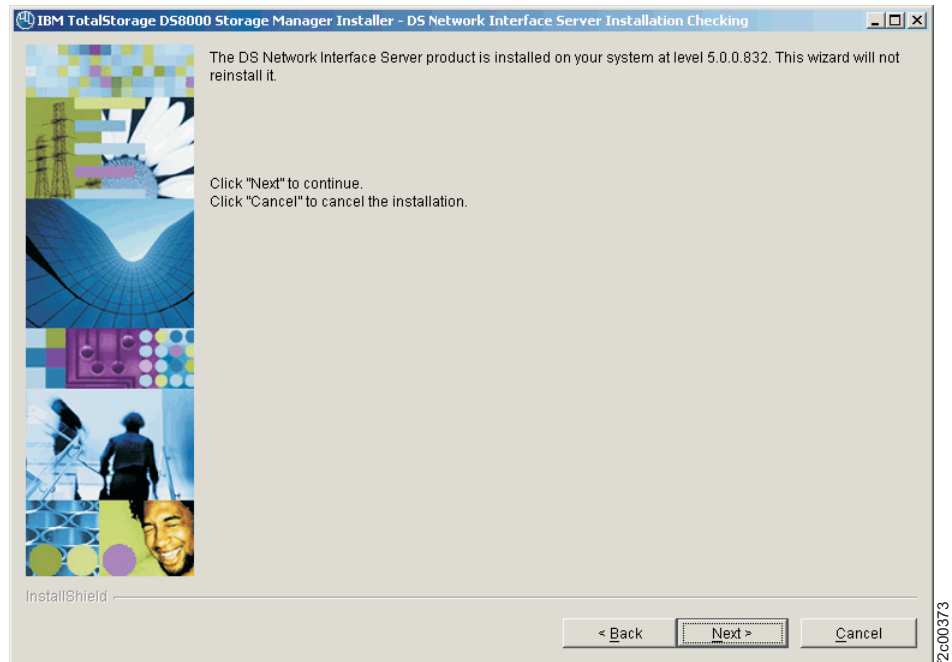
11. Observe the information that is displayed on the Linux Product Version Check window regarding the version of the product that is already installed on your

system and the action the wizard will perform. Click **Next** to perform the action or click **Cancel** to exit the upgrade. When you click **Next**, the DS Storage Manager Server Installation Checking window is displayed if the DS Storage Manager service is installed on your system.



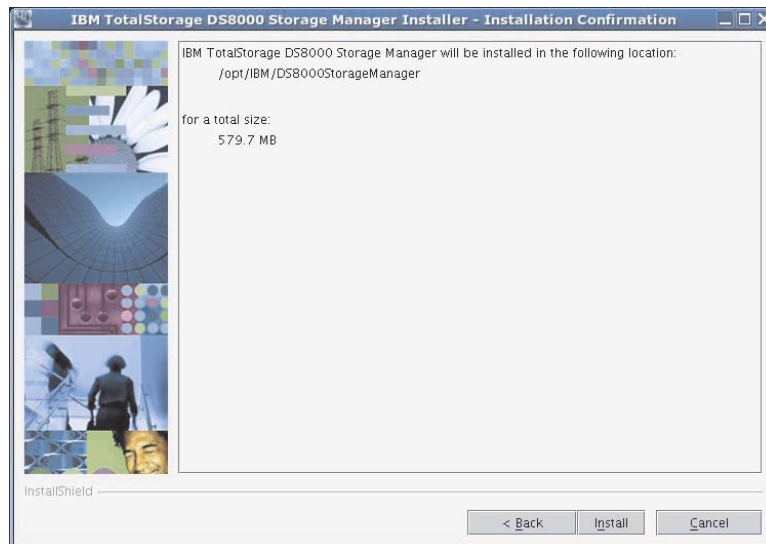
DS Storage Manager Server Installation Checking window

12. Observe the information that is displayed on the Linux DS Storage Manager Server Installation Checking window regarding the version of the DS Storage Manager Server that is already installed on your system and the action the wizard will perform. If you agree, click **Next** or click **Cancel** to exit the upgrade process. When you click **Next**, the DS Network Interface Server Installation Checking window is displayed if the DS Network Interface service is installed on your system.



Linux DS Network Interface Server Installation Checking window

13. Observe the information that is displayed on the Linux DS Network Interface Server Installation Checking window regarding the version of the DS Network Interface Server that is already installed on your system and the action the wizard will perform. If you agree, click **Next** or click **Cancel** to exit the upgrade process. When you click **Next**, the Installation Confirmation window is displayed. This window displays both the location where the product will be installed and the total size that is required for the upgrade.



Linux Installation Confirmation window

14. Click the **Install** button on the Installation Confirmation window to begin the upgrade process. There are several progress windows that are displayed. There is no required interaction on your part for each of the progress windows that are displayed. However, you can choose to cancel (not recommended) the installation on any of the progress windows with varying consequences. The installation process performs the following actions in sequence:

- a. If the two servers (DS Storage Manager Server and DS Network Interface Server) are already installed on your system, they are stopped in the following order of windows:
 - 1) The Embedded IBM WebSphere Application Server - Express server (part of the DS Storage Manager Server) window is stopped first.
 - 2) The service window (DS Network Interface Server) is stopped next.
 - 3) The WS Help System (part of the DS Storage Manager Server) window is stopped, if it was not stopped before by the Embedded IBM WebSphere Application Server.
- b. If one or both of the servers are not installed or have to be upgraded on your system, they are installed or upgraded in the following order (the progress of the installation is indicated on the associated progress window):
 - 1) DS Storage Manager Server Installation Progress window
 - 2) DS Network Interface Server Installation Progress window

Note: You can click **Cancel** (not recommended) during the DS Network Interface Server installation process. The process does not stop immediately when the **Cancel** button is clicked. Rather, the process continues to install all the files that are associated with this part of the installation. These files remain installed and are not reinstalled upon reactivation of the installation process.

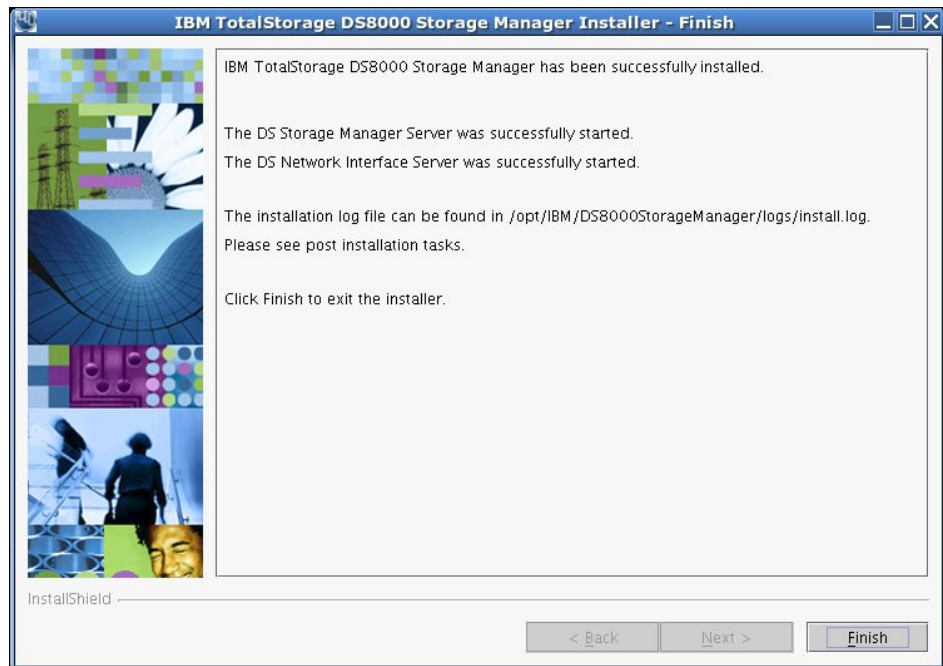
A confirmation message is displayed asking you to confirm that you want to cancel the installation.

- c. The Components Installation Progress (displaying the installation or upgrade progress of the DS Storage Manager product applications) is displayed after the servers have been installed or upgraded.

Note: You can click **Cancel** (not recommended) during the components installation process. The installation process stops immediately when the **Cancel** button is clicked. A window with a confirmation message is displayed. For a new installation, when you confirm that you want to leave, all the files that have been copied up to the point that the **Cancel** button was clicked are uninstalled. You are then exited from the installation process.

When this part of the upgrade is completed, the system starts both servers: first the DS Network Interface Server and then the DS Storage Manager Server. Wait for the servers to be started before going to the next step. When the servers have been started, the Installer Finish window is displayed.

15. Click **Finish** to exit the upgrade process. When the installation wizard exits, a file (postinstallation.txt) is displayed, containing the postinstallation tasks. Follow the instructions in the postinstallation text file to complete the postinstallation tasks. If the upgrade failed, you must exit the upgrade process and check the install log for error messages.



Linux DS8000 Storage Manager Installer Finish window

16. If the Finish window indicates that the upgrade has failed, check the install log for error messages. The install log is located in *xxx/logs/install.log*, where *xxx* is the destination directory where the IBM TotalStorage DS Storage Manager is installed.
Errors fall into two categories: system or incorrect values that are designated during the install process. Use the following guidelines to correct these errors:
 - If the error is due to a system problem, correct it and reinstall the GUI using either the graphical or silent mode of upgrade.
 - If the error is due to a wrong upgrade parameter value, restart the upgrade using the steps that are described in this procedure or the steps that are described in the silent mode upgrade. Navigate to the Console Server Parameters window and insert the correct values. Then finish the upgrade.
17. Check the postinstallation instructions after the installation has completed. The *postinstallation.txt* file is displayed immediately after the installation of the DS Storage Manager finishes processing. If the *postallation.txt* file does not display, click the **Post installation tasks** link on the LaunchPad window to access these instructions. These instructions provide reminders and the steps to help you start using the DS Storage Manager.
18. Exit the LaunchPad program by clicking **Exit** on the LaunchPad window.
19. Type the following command when you are finished using the IBM TotalStorage DS Storage Manager CD and you want to remove the CD from your computer:
umount /mnt/cdrom where */mnt/cdrom* is where the CD is mounted.

Upgrading the IBM TotalStorage DS8000 Storage Manager on the Linux operating system in unattended (silent) mode

You can upgrade the IBM TotalStorage DS Storage Manager in your Linux environment using the unattended (silent) mode.

The unattended (silent mode) upgrade option allows you to run the upgrade program unattended. Use this method of upgrade to issue a command from a command prompt window from the root directory of the IBM TotalStorage DS Storage Manager CD.

Perform the following steps to upgrade the IBM TotalStorage DS Storage Manager in your Linux environment using the unattended mode:

1. Log on as a user with root authority.
2. Create a mount point or choose an existing mount point.
3. Determine where you want to mount the DS8000 Storage Manager CD. For example, imagine you have decided to mount the CD in a CD-ROM drive known as /dev/cdrom at mount point /mnt/cdrom. When the time comes for you to mount the CD in this drive, perform the following steps:
 - a. Insert the CD in the CD-ROM drive.
 - b. Type the following command at the command prompt: **# mount -o ro /dev/cdrom /mnt/cdrom**
4. Change the current directory to the mount point for the CD-ROM drive. For example, if the CD was mounted at the /mnt/cdrom mount point, type the following command: **# cd /mnt/cdrom**
5. Read the README.txt file that is located on the DS8000 Storage Manager CD to check for information that might supersede the information in these upgrade instructions.
6. Type **./setup.linux -silent** in the same Command Prompt window in which you changed the current directory to the root of the CD, to start the DS8000 Storage Manager upgrade program in silent mode.
7. Wait for the DS8000 Storage Manager installation program to install the product. This can take 5 - 10 minutes.
8. Review the log file for any possible error messages, after the installation program completes. If there are no errors, the installation program completed successfully. The log file is located in xxx/logs/install.log, where xxx is the destination directory where the DS8000 Storage Manager was installed. The install.log file contains a trace of the installation actions.
9. Perform the postinstallation tasks after the installation has completed. You can read the instructions for the postinstallation tasks in the postinstallation.txt in the doc directory on the product CD.

Chapter 4. DS Storage Manager postinstallation instructions

You have just installed the DS Storage Manager and its two supporting servers. The servers are automatically started when you start your computer and remain on until you stop them manually or there is a system outage. The following instructions describe how to log on to DS Storage Manager and how to begin using it.

Additional instructions are provided for checking the status of the DS Storage Manager servers, how to stop and start them, and how to stop and start the DS Storage Manager itself in the Windows and Linux operating systems.

Internet browser support

The DS Storage Manager can be used on the Internet Explorer (IE) and Netscape Navigator versions of Internet browsers.

The following list shows the Internet browser versions that support the use of the DS Storage Manager.

- IE 6.x
- Netscape 6.2
- Netscape 7.x

Note: Unless otherwise noted, these browsers work with the DS Storage Manager installed on the Windows operating system.

Stop and start the DS Storage Manager Servers

The IBM TotalStorage DS Storage Manager server and IBM TotalStorage DS Network server are installed and activated when you installed the DS Storage Manager. These servers remain active until you stop them or there is a system failure.

Each server is accessed by a different method depending on your operating system.

Stopping and restarting the DS Storage Manager servers on a Windows operating system

You can stop or restart the DS Storage Manager servers by using the Windows Programs list.

Log on to your Windows operating system and use the following steps to work with the DS Storage Manager servers. Perform these steps for each server.

1. Click **Start**.
2. Select **Programs** to display the programs list.
3. Click the server (IBM TotalStorage DS Storage Manager or IBM TotalStorage DS Network Interface) that you want to stop or restart.
4. Click **Stop**, or **Start** for the action that you want to complete.

Stopping and restarting the DS Storage Manager servers on a Linux operating system

You can type the following commands to check on the status of the servers and then to stop or restart them on a Linux operating system.

Open a command prompt window and use the following commands to check status and stop or start the servers.

1. Type `/etc/init.d/dsStorageManagerServer status` to see the status of the DS Storage Manager server.
2. Type `/etc/init.d/dsStorageManagerServer stop` to stop the DS Storage Manager server.
3. Type `/etc/init.d/dsStorageManagerServer start` to start the DS Storage Manager server.
4. Type `/etc/init.d/dsnisrv status` to see the status of the DS Network Interface server.
5. Type `/etc/init.d/dsnisrv stop` to stop the DS Network Interface server.
6. Type `/etc/init.d/dsnisrv start` to start the DS Network Interface server.

Starting the DS Storage Manager program

The topics in this section describe how to start the DS Storage Manager program.

The DS Storage Manager program times out over a period of inactivity or when you turn off your computer. The way you start the DS Storage Manager is different depending on whether you have a Windows or Linux operating system.

Starting the DS8000 Storage Manager on a Windows operating system

You can open the DS8000 Storage Manager using the Programs list. The DS8000 Storage Manager opens in your default browser.

Log on to your Windows operating system and use the following steps to access the DS8000 Storage Manager program.

1. Click **Start**.
2. Click **Programs**.
3. Click **IBM TotalStorage Manager DS8000 Storage Manager** and then click **Open DS Storage Manager**. The DS8000 Storage Manager is opened in the browser that you have set as your default.

When you start the DS8000 Storage Manager, the IBM TotalStorage DS8000 Signon window is displayed. Enter the user name and password to access the program. The default user name is **admin**, and the default password is **admin**.

Starting the DS8000 Storage Manager on a Linux operating system

You can start the DS8000 Storage Manager on a Linux operating system by issuing a command from the command prompt window. Within this command, you must specify the browser that you are using.

Log on to your Linux operating system and follow these steps to start the DS8000 Storage Manager program:

1. Open a command prompt window for Linux.
2. Type `browsername https://<host-name>:<port-value> /DS8000/Console`

browsername

This is the name of the browser that you used to open the DS8000 Storage Manager console.

<host-name>

This is the fully qualified host name of the machine where the DS Storage Manager server is installed.

<port-value>

This is the HTTPS port number (the default value is 8452) that the DS Storage Manager server uses for a secure HTTP transport. The port number must be a port that is not used by another process on the system.

When you start the DS8000 Storage Manager console, the IBM TotalStorage DS8000 Signon window is displayed. Enter your user name and password to access the program.

Getting started with the DS Storage Manager console

Configuration capabilities and the use of the Copy Services features after the initial installation is dependant on your purchase.

The following components are available to you within the DS8000 Storage Manager program:

- The simulated configuration component
- The real-time configuration component with Copy Services

All of the functions that are associated with these components are not described here. If you need help with any of the pages that you open, you can click the question mark icon on the tool bar. Page help is displayed in another window along with access to an Information Center that contains information about the DS8000.

Consider the following items as first steps in the use of either of these components.

Log In

You can log in to the DS8000 Storage Manager by providing your user name and password. This function is generally administered through your system administrator and by your company policies.

Note:

- The default password that you use the first time (generally the password is **adminadmin**) expires immediately. You must use the **Change password** feature to ensure your further use of the GUI.
- Choose a password that you can readily remember because the password is not available to your administrator nor is it retrievable from the system.

Simulated component

This component is available if you have installed the DS8000 Storage Manager in the offline mode (your 2107 is preinstalled with the real-time component of the DS8000 Storage Manager). You can begin using the simulated component immediately after logging into the DS Storage Manager. This component provides the ability to create or modify logical configurations when your storage unit is disconnected from the network. After creating the configuration, you can save it and then apply it to a network-attached DS8000 at a later time.

Activate your licenses

Before any configuration can be applied to your DS8000 network, you must activate your license activation codes according to the processes that are described for activating your licenses.

Real-time component (preinstalled on 2107)

Use the real-time component selections of the DS8000 Storage Manager to work with the logical configuration and Copy Services features when your storage unit is attached to a DS8000 network.

Copy Services

You can use the Copy Services selections of the Real-time Manager if you have purchased the optional Copy Services features.

Activating licenses

After the IBM service representatives have completed your DS8000 storage complex installation, your first step is to activate your licenses.

To activate your licenses, you must perform the following actions:

- Obtain your feature activation codes.
- Apply the activation codes to your storage unit. You can apply the activation codes by importing a file that you download from the IBM Disk Storage Feature Activation (DSFA) Web site.

The initial enablement of any optional DS8000 licensed function is a concurrent activity (assuming the appropriate level of microcode is installed on the machine for the given function).

Note: The following activating activities are disruptive and require a machine IML (Model 921 and Model 922) or reboot of the affected image (Model 9A2):

- Removal of a DS8000 licensed function to deactivate the function.
- A lateral change or reduction in the license scope. A lateral change is defined as changing the license scope from fixed block (FB) to count key data (CKD) or from CKD to FB. A reduction is defined as changing the license scope from all physical capacity (ALL) to only FB or only CKD capacity.

Obtaining activation codes

To obtain your feature activation codes for the licensed features for each storage unit, you must connect to the IBM Disk Storage Feature Activation (DSFA) Web site.

Before connecting to the site, ensure that you have the following items:

- The IBM License Function Authorization documents. If you are activating codes for a new storage unit, these documents are included in the shipment of the storage unit. If you are activating codes for an existing storage unit, IBM sends these documents to you in an envelope.
 - A diskette for downloading your activation codes into a file if you cannot access the IBM TotalStorage DS Storage Manager from the system that you are using to access the DSFA Web site. Instead of using a diskette, you can also write down the activation codes and then go over to the system that runs the DS Storage Manager and manually enter them.
1. Start the DS Storage Manager application.
 2. In the navigation panel select **Real-time ManagerManage Hardware**, then select **Storage Units**.

3. In the storage unit main page select the storage unit, click **Properties** in the **Select Action** drop-down list, and then click **Go**. The properties page displays for the storage unit.
4. Click the **General** tab, and the general properties display in the panel. Gather the following information about your storage unit. You must enter this information at the IBM Web site in the next step of this task. You can use the following table to document this information.
 - From the **MTMS** field, note the machine's serial number, and the model number. The Machine Type - Model Number - Serial Number (MTMS) is a string that contains the machine type, model number, and serial number. Only the last seven characters of the string are the machine's serial number.
 - From the **Machine signature** field, note the machine signature.

Property	Your Storage Unit's Information
Machine's serial number	
Model number	
Machine signature	

5. At a computer with an Internet connection and a browser, connect to the IBM Disk Storage Feature Activation (DSFA) Web site at <http://www.ibm.com/storage/dsfa>.
6. The DSFA application displays in the browser. Use the application to obtain the activation codes and follow the instructions on the screen.

Note: In most situations, the DSFA application can locate your 2244 license authorization record when you enter the DS8000 (2107) serial number and signature. However, if the 2244 license authorization record is not attached to the 2107 record, you must assign it to the 2107 record in the DSFA application. In this situation, you will need the 2244 serial number (which you can find on the License Function Authorization document).

Applying activation codes

Use this process to apply the activation codes that enable you to begin configuring storage on a storage image.

You cannot have both the Apply activation codes page and the Import activation codes page open at the same time. You must close one in order to access the other.

Note: Before you begin this task, you must resolve any current DS8000 problems. Contact IBM Support for assistance in resolving these problems.

The easiest way to apply the feature activation codes is to download the activation codes from the IBM Disk Storage Feature Activation (DSFA) Web site to your local computer and then to import the file into the DS Storage Manager. If you cannot access the DS Storage Manager from the same computer that you used to access the DSFA Web site, you can download the file to a diskette or write down the information. If you are using either of these latter methods, ensure that you have your diskette containing the downloaded activation codes file or your paper that contains the written activation codes before you begin the following steps.

1. In the navigation panel, select, in order, **Real-time Manager**, **Manage Hardware**, and **Storage images**. In Storage images — Main Page, select a Storage image. Then select **Apply activation codes** in the **Select Action** drop-down list. Then click **Go**. The Apply activation codes page is displayed.

- a. If you already imported your activation codes from a file or retrieved existing codes from the storage unit, the values are displayed in the fields and you can modify or overwrite them, as appropriate.
- b. If you are importing your activation codes from a file that you downloaded from the DSFA Web site, click **Import key file**. Once you complete the import process, the data from the file is displayed.
- c. If you did not download your activation codes into a file, enter the codes into the appropriate fields.

Note: For model 9A2, you need to perform this process for both images, one image at a time.

Note: The **Capacity** and **Storage type** fields are populated based on the information contained within the activation codes.

2. Click **Apply**, then **OK** to complete the process.

Importing activation codes

Use this process to import the activation codes that must be applied before you can begin configuring storage on a storage image.

Note: Before you begin this task, you must resolve any current DS8000 problems. Contact IBM Support for assistance in resolving these problems.

1. In the navigation panel, under Manage Hardware, select Storage images. In Storage images — Main Page, select a Storage image. Then select **Apply activation codes** in the **Select Action** drop-down list. Then click **Go**. The Apply activation codes page is displayed.
2. Click **Import key file**. The Import page is displayed.
3. In the **Select file to import** field, specify the target file. Use the **Browse** button to navigate to the appropriate directory.
4. After you have specified the code file, click **OK** to complete the process.

Chapter 5. IBM TotalStorage Management Console

The IBM TotalStorage Management Console is the focal point for configuration, copy services management, and maintenance activities.

The Management Console is a dedicated workstation that is physically located (installed) inside your DS8100 and DS8300 and can automatically monitor the state of your system, notifying you and IBM when service is required. The Management Console can also be connected to your network to enable centralized management of your system using the IBM TotalStorage DS Command-Line Interface or the storage management software that uses the IBM TotalStorage DS Open API.

To provide continuous availability of customer access to the management console functions, having a second management console is recommended. This second management console can be provided in two ways:

- **External** (outside the 2107) - This console is generally installed in the customer-provided rack. It uses the same hardware as the internal management console.
- **Internal** (a second management console that resides within a second 2107) - This console can be cabled together with the primary internal management console of the first 2107.'

Logging onto the Management Console

When logging onto the Management Console, the customer initially uses the management console predefined user ID (**customer**) and password (**cust0mer**). After logging in the first time, the customer can choose to change the password.

The following log on process provides access to the WebSM interface. From this interface you can change your password or access the DS Storage Manager.

Note: CEs use a different ID and password when using the Management Console.

1. Slide the Management Console keyboard/display drawer out from the machine and raise the panel display.
2. Turn on the console. The login window is displayed, prompting you to provide your user ID and password. The WebSM interface window is displayed.
3. Choose what you want to do; access the DS Storage Manager or change your password.

Note: You cannot use the predefined password after you change it. If you forget your password, you must contact your IBM service representative to have the password reset.

Logging off the Management Console

The management console must always be left operational. However, you might consider the following actions when you leave the management console unattended or when you must logoff your management console session.

1. In the main menu, click **Console** —> **Exit**. At this point, you can choose to save the state of the console for the next session by clicking the box next to the option.
2. Click **Exit Now**.

3. When you exit from management console session, you can choose to shut down, reboot, or log off your session. The following is a description of each option:

Logout

Returns the user to the login prompt without shutting down the Management Console. This is an appropriate action when you leave the management console unattended.

Reboot Console

Shuts down the management console and then reboots it to the login prompt. This might be necessary when your management console becomes hung up and a service provider is not available to pursue the cause of this event.

Shutdown Console

Powers off the management console. Generally, this action is done when there is an emergency or when the service provider is repairing management console.

Changing the Management Console password

You can change the password you use to log on to the management console. When the password is changed, the predefined password can no longer be used.

Consider the following factors when you make a change to the management console password:

- The new password must be at least seven characters and contain a numeric character.
 - The numeric character cannot be in the first or last position of the password.
 - The new password nullifies the use of the predefined password. Thus, if you forget your new password, you must contact your IBM service representative to have the password reset.
 - Passwords are case-sensitive.
 - After inputting your new password, ensure that you record it for your personal retrieval.
1. Select **Change HMC password**. The change user password window is displayed with three input fields.
 2. Type your old (current) password into the first input field.
 3. Type your new password into the New password and Confirm password input fields.
 4. Click **OK** and the system accepts the changed password. The next time you log onto the management console you can use your new password.

Accessing the DS Storage Manager

You can access the DS Storage Manager after the 2107 installation using the Management Console or you can use your PC. It is important that you connect to the DS Storage Manager using one of the supported Internet browser systems. When you use the Management Console, log on to the WebSM interface and from here you can access the DS Storage Manager using the internet browser system. Use the following process to connect to the DS Storage Manager using the Management Console.

1. Log on to the WebSM interface at the Management Console using the predefined user ID and password.

2. Right click on the desktop until the Fluxbox - 0.9.4 is displayed.
3. Click the **Net** selection in the Fluxbox. Another box is displayed with the words Net and Browser.
4. Click **Browser**. The Web Browser is started with no address bar and a Web page titled WELCOME TO THE DS8000MANAGEMENT CONSOLE is displayed.
5. Click the link "IBM TotalStorage DS Storage Manager" to access the DS Storage Manager.

Chapter 6. Administering user accounts

The topics in this section provide information related to administering your DS8000 user accounts. If you have administrator level privileges, you can add a new user account, delete an existing user account, or modify the user name, password, and group for a user account. There might be times when users forget the password that they use to access the DS Storage Manager. After going beyond the set number of allowable attempts with the wrong password, the account is locked. The administrator can unlock a user account. If the Admin account is locked, the administrator must use the security recovery utility tool. If you do not have administrator level privileges, your account administration privileges are limited to entering a new password for your own user ID.

Adding user accounts

You must have Administrator level privileges to add a new user account. If you do not have Administrator level privileges, Add does not display in the **Select Action** drop-down box of the User administration - Main page. You can also add a user account with the **mkuser** DS Command-Line Interface command. See DS CLI documentation for more information.

1. Under **Monitor system**, select **User administration**. In User administration — Main page, select **Add** in the **Select Action** drop-down box. Then click **Go**. The Add user page is displayed.
2. Specify the user name. You can enter up to 16 characters.
3. Specify the user account password. This entry is displayed as asterisks.
Passwords must meet the following criteria:
 - Passwords must contain at least 5 alphabetic characters, and at least one numeric character, with an alphabetic character in the first and last positions.
 - Passwords are limited to a total of 16 characters.
 - Passwords cannot contain the user name.
 - Four unique new passwords must be issued before an old password can be reused.
4. Retype the password in the **Confirm password** box. This entry must match the password entry above. Characters in this field appear as asterisks.
5. Assign a group role to the user account by selecting the appropriate group in the **Group assignment** boxes.
6. Click **Ok**. The user ID that you added is available for selection in the User administration — Main Page.

Modifying user accounts

If you have Administrator level privileges, you can modify the user name, password, and group for a user account. If you do not have Administrator level privileges, you can only enter a new password for your own user ID. You can also modify a user account with the **chuser** DS Command-Line Interface command. See DS CLI documentation for more information.

1. Under **Monitor system**, select **User administration**. In User administration — Main page, select **Modify** in the **Select Action** drop-down box. Then click **Go**. The Modify user page is displayed.
2. To modify the user name, enter up to 16 characters. If you do not have Administrator level privileges, your user name appears in the **User name** box by default and you are not able to modify it.

3. Specify the user account password. Passwords must contain at least 5 alphabetic characters, and at least one numeric character, with an alphabetic character in the first and last positions. Passwords are limited to a total of 16 characters. The user name can not be part of the password. The minimum number of unique new passwords that must be used before an old password can be reused is four. This entry will appear as asterisks.
4. Retype the password in the **Confirm password** box. This entry must match the password entry above. Characters in this field appear as asterisks.
5. Assign a group role to the user account by selecting the appropriate group in the **Group assignment** box.
6. Click **Ok**. The properties for the user account are immediately modified.

Removing user accounts

If you have Administrator level privileges, you can remove an existing user account. If you do not have Administrator level privileges, Delete does not display in the **Select Action** drop-down box of the User administration — Main page. You can also remove a user account with the **rmuser** DS Command-Line Interface command. See DS CLI documentation for more information.

1. Under **Monitor system**, select **User administration**. In User administration — Main page, select the user ID to remove.
2. Select **Delete** in the **Select Action** drop-down box. Then click **Go**. A confirmation message is displayed.
3. Click **Ok**. The user ID is removed immediately.

Unlocking a user account

There might be times when users forget the password that they use to access the DS Storage Manager. After going beyond the set number of allowable attempts with the wrong password, the account is locked. To unlock a user's account, the administrator can use the unlock user process. If the Admin account is locked, the Administrator must use the security recovery utility tool. You can also unlock a user account with the **chuser** DS Command-Line Interface command. See DS CLI documentation for more information.

You must have administrator-level privileges to unlock a user's account. If you do not have administrator-level privileges, you cannot unlock a user account or use the security recovery tool to unlock the Admin account.

Note: This task only explains how to use the unlock user process. The Unlocking an administrative password task describes how to use the security recovery utility tool to unlock the Admin account.

1. Under **Monitor system**, select **User administration**. In User administration — Main page, select the user ID to unlock.
2. Select **Unlock user** in the **Select Action** drop-down box. Then click **Go**. A confirmation message is displayed.
3. Click **OK**. The Account Status column updates accordingly.

Unlocking an administrative password

There might be times when administrative users forget the password that they use to access the DS Storage Manager. After going beyond the set number of allowable attempts with the wrong password, the account is locked. If the Admin account is involved, the administrator must use the security recovery utility tool. You cannot unlock an administrative password using the DS Command-Line Interface. The administrative user is forced to establish a new password. Using the **chuser** command, you can specify a password that expires after the initial use, then create a new password. See DS CLI documentation for more information.

Note:

- This task only explains how to use the security recovery utility tool to unlock the Admin account. The topic "Unlocking a user account" describes how to unlock a nonadministrative user account.
 - The security recovery utility tool is a script that is installed in a file directory. You run the script from the directory.
1. Access the directory where the recovery tool (script) has been installed. The script resides in the following directories, depending on where the DS Storage Manager is installed:
 - On an HMC, the directory is: `/opt/essni/bin/`
 - On a Windows operating system, the directory is: `C:\Program Files\IBM\dsniserver\bin\`
 2. Type the script name. The script name depends on where the DS Storage Manager is installed:
 - On an HMC, the script name is: `securityRecoveryUtility.sh -r`
 - On a Windows operating system, the script name is:
`securityRecoveryUtility.bat -r`
 3. Press the **Enter** key. The script runs and the Admin account is unlocked.

Chapter 7. Storage system logical partition

A storage system logical partition allows the division of a single server into two completely independent virtual storage systems.

A logical partition uses hardware and firmware to logically allocate the resources between partitions.

The storage system logical partition function is currently supported on the DS8300 Model 9A2. It provides two virtual storage systems in one physical machine. Each storage system logical partition can run its own level of licensed machine code.

The resource allocation for processors, memory, and I/O slots in the two storage system LPARs on the DS8300 is divided into a fixed ratio of 50/50.

Note:

- A storage system logical partition must be associated with the same storage system logical partition.
- Each storage system logical partition can have 255 logical subsystems (LSSs) with LSS IDs 0–254.

DS8300 LPAR implementation

The DS8300 Model 9A2 provides a split between the resources in a 50/50 ratio.

Each storage image uses the machine type/model number/serial number of the DS8300 Model 9A2 base frame. The frame serial number ends with a 0. The last character of the serial number is replaced by a number in the range one to eight that uniquely identifies the DS8000 image. This character is a 1 or a 2, because there are only two available storage images. The serial number distinguishes between the storage images in the GUI and CLI. The serial number is also used for licensing and the allocation of licenses between the storage images.

The following figure illustrates the DS8300 LPAR resource allocation that is allowed.

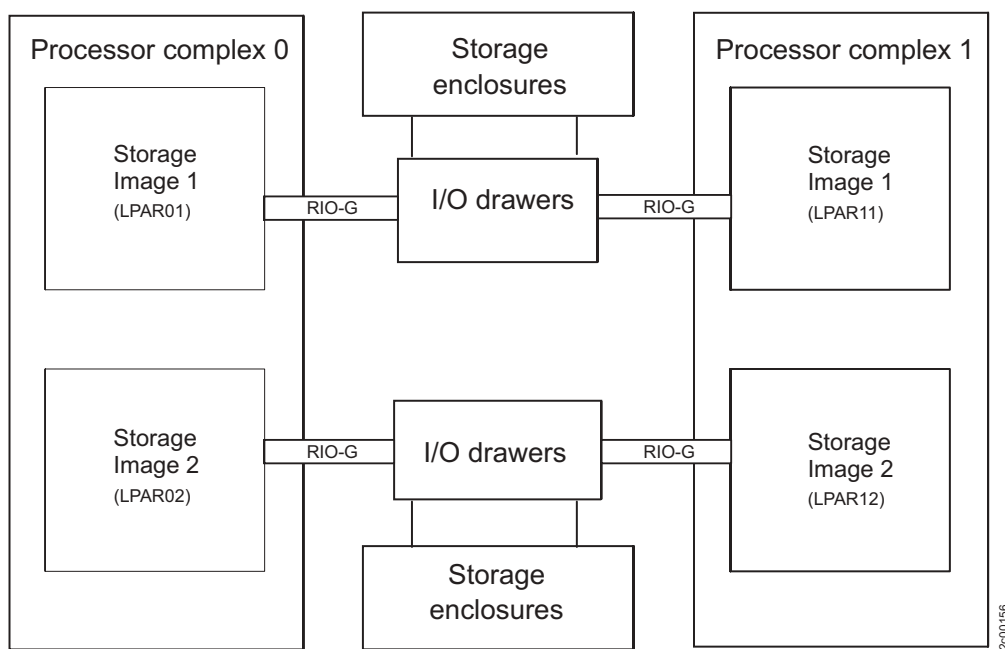


Figure 6. DS8300 LPAR resource allocation

Each storage image has access to the following components:

- 50 percent of the processors
- 50 percent of the processor memory
- Up to 16 host adapters (4 I/O drawers with up to 4 host adapters)
- Up to 320 disk drives (up to 96 TB of capacity)

Required hardware components to build a storage image

When you build a storage image, you need to consider your I/O resources, the processor and memory allocations, and the RIO-G interconnect separation.

The Hardware Management Console (HMC) manages the resource allocation on a pSeries between LPARs. However, the DS8300 comes preconfigured with all required LPAR resources assigned to either storage image.

Figure 7 on page 97 shows the split of all available resources between the two storage images. Each storage image has 50% of all available resources.

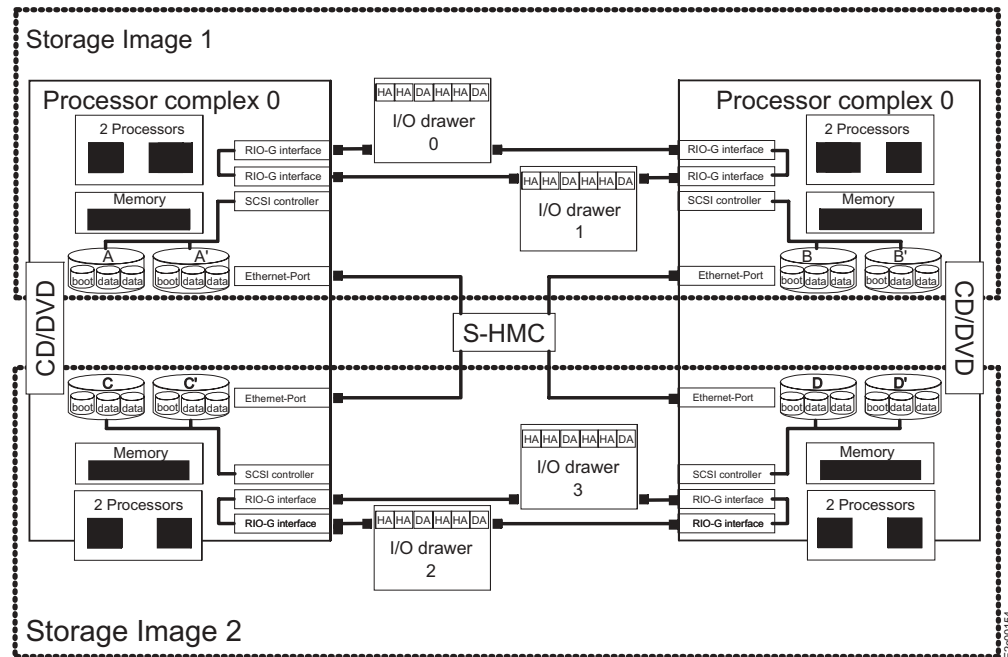


Figure 7. Storage image resource allocation in the processor complexes of the DS8300

I/O resources

The I/O resources that are required for a storage image consist of various components.

For one storage image, the following hardware resources are required:

- 2 SCSI controllers with 2 disk drives each
- 2 Ethernet ports (to communicate with the HMC)
- 1 Thin Device Media Bay (for example, a CD or a DVD which can be shared between the LPARs)

Each storage image has two physical disk drives in each processor complex. Each disk drive contains three logical volumes, the boot volume and two logical volumes for the memory save dump function. These three logical volumes are mirrored across the two physical disk drives for each LPAR. For the DS8300 Model 9A2, there are four drives in one physical processor complex.

Processor and memory allocation

In the DS8300 Model 9A2 each processor complex has four processors and up to 128 GB memory capacity. There is also a 50/50 split for processor and memory allocation. Therefore, every LPAR has two processors and every storage image has four processors.

The memory limit depends on the total amount of available memory in the whole system. The following memory allocations per storage image are available:

- 32 GB (16 GB per processor complex, 16 GB per storage image)
- 64 GB (32 GB per processor complex, 32 GB per storage image)
- 128 GB (64 GB per processor complex, 64 GB per storage image)

- 256 GB (128 GB per processor complex, 128 GB per storage image)

RIO-G interconnect separation

Each RIO-G loop is dedicated to a given storage image. All I/O enclosures on the RIO-G loop with the associated host adapters and drive adapters are dedicated to the storage image that owns the RIO-G loop.

As a result of the strict separation of the two images, the following configuration options exist:

- Each storage image is assigned to one dedicated RIO-G loop; if an image is offline, its RIO-G loop is not available.
- All I/O enclosures on a given RIO-G loop are dedicated to the image that owns the RIO-G loop.
- Host adapter and device adapters on a given loop are dedicated to the associated image that owns this RIO-G loop.
- Storage enclosures and storage devices behind a given device adapter pair are dedicated to the image that owns the RIO-G loop.
- Configuring capacity to an image is managed through the placement of storage enclosures on a specific DA pair that is dedicated to this image.

Chapter 8. Optional feature codes for the DS8000

The DS8000 is designed to meet or exceed the storage requirements of a broad range of operating environments.

There are many features that are enabled to meet your data management requirements. The following two areas are particularly useful for data management.

- Parallel Access Volumes (PAV) used for zSeries environments
- Advanced Copy Services functions

Parallel Access Volumes

This feature allows multiple z/OS jobs to execute in parallel and queue management on the DS8000. Dynamic PAVs autonomically reconfigure the storage subsystem that is based on your workload.

Advanced Copy Services functions

The advanced Copy Services functions are sometimes referred to as the IBM TotalStorage Resiliency Family. These Copy Services functions offer a variety of data duplication, data migration, and disaster recovery solutions that are supported on open systems, zSeries®, and Tivoli® Storage Manager.

The following features are key components of the IBM TotalStorage Resiliency Core Technology:

- IBM TotalStorage Point-in-time copy, previously known as FlashCopy®
- Remote Mirror and Copy, previously known as Peer-to-Peer Remote Mirror (PPRC)
 - IBM TotalStorage Metro Mirror, previously known as Synchronous PPRC
 - IBM TotalStorage Global Mirror, previously known as Asynchronous PPRC
 - IBM TotalStorage Global Copy, previously known as PPRC Extended Distance
- IBM TotalStorage z/OS® Global Mirror, (previously known as Extended Remote Copy [XRC])

Note: z/OS Global Mirror (previously known as Extended Remote Copy or XRC) provides a long-distance remote copy solution across two sites for open systems and z/OS data using asynchronous technology. It is based on the existing Copy Services functions called Global Copy and FlashCopy and additional new functions that form consistency groups across one or more storage units to manage and control Global Mirror sessions.

Using parallel access volumes for S/390 and zSeries hosts

This section provides the following guidelines for using parallel access volumes (PAVs):

- An overview of PAVs
- PAV requirements
- An overview of the input/output configuration program (IOCP)
- Guidelines for defining the logical control units
- Guidelines for reconfiguring the logical control units

Parallel access volumes

Parallel access volumes (PAVs) allow your system to access volumes in parallel.

A PAV capability represents a significant performance improvement by the storage unit over traditional I/O processing. With PAVs, your system can access a single volume from a single host with multiple concurrent requests.

You must configure both your storage unit and operating system to use PAVs. You can use the logical configuration definition to define PAV-bases, PAV-aliases, and their relationship in the storage unit hardware. This unit address relationship creates a single logical volume, allowing concurrent I/O operations.

Static PAVs are PAVs in which the association between a PAV-base address and its PAV-aliases is pre-defined and fixed. That is, the PAV-aliases of a PAV-base address remain unchanged. Dynamic PAVs, on the other hand, are PAVs in which the association between a PAV-base address and its PAV-aliases is dynamic. The device number types (PAV-alias or PAV-base) must match the unit address types as defined in the storage unit hardware.

PAVs can improve the performance of large volumes. You get better performance with one base and two aliases on a 3390 Model 9 than from three 3390 Model 3 volumes with no PAV support. With one base, it also reduces storage management costs that are associated with maintaining large numbers of volumes. The alias provides an alternate path to the base device. For example, a 3380 or a 3390 with one alias has only one device to write to, but can use two paths.

The storage unit supports concurrent or parallel data transfer operations to or from the same volume from the same system or system image for S/390 or zSeries hosts. An S/390 with PAV software support enables multiple users and jobs to simultaneously access a logical volume. Read and write operations can be accessed simultaneously to different domains. (The domain of an I/O operation is the specified extents to which the I/O operation applies.)

PAV requirements for the DS8000

Parallel access volumes (PAVs) can provide significant performance enhancements in S/390 and zSeries environments by enabling simultaneous processing for multiple I/O operations to the same logical volume.

The following list identifies the requirements for using PAVs on the DS8000:

- The configuration that you create on the DS8000 must match the configuration that you create using IOCP or HCD.
- The control unit (CU) image numbers that exist on the DS8000 depend on what logical subsystems (LSSs) you configured:
 - The CU image addresses are 0 - 254. For IOCP and HCD, the CU addresses are hex 00 - FE.
 - Logical subsystem (LSS) addresses are 0 - 254.

The configuration might only have a subset of these addresses, and the subset might not be contiguous.

You need to know what CU images you created before you can update the I/O configuration. Your host system uses the updated I/O information in order to know what CU images are on the DS8000.

Overview of the input/output configuration program for the DS8000

The input/output configuration program provides the information needed for the DS8000 when handling PAVs.

The output that the input/output configuration program (IOCP) generates, contains the I/O definition for the channel subsystem of an S/390 or zSeries host. It is not used as input to the I/O definition for the OS/390®, z/OS®, VM, z/VM, or VSE operating system. Therefore, you do not need to distinguish between bases and alias devices in the input to IOCP. However, doing so provides an accurate representation of the I/O configuration.

The OS/390 and z/OS operating systems require that the input/output definition file (IODF) that the hardware configuration definition (HCD) creates, specifies base and alias devices. The OS/390 and z/OS operating system dynamically discover which base and alias devices are associated with each other. HCD invokes IOCP to build an IOCDS that identifies the I/O definition for the host processor.

The VM, z/VM, and VSE operating systems dynamically discover which devices are bases and aliases and which are associated with each other. The VM and z/VM operating systems support PAVs for guest usage.

An IOCP example

The following is an IOCP example that defines ESCON and FICON bridge channel path IDs that you have attached to an ESCON director. This example uses the storage unit machine number, 2107. HCD users can use the information in the following IOCP example to determine how to define a storage unit to HCD.

```
*****
*                                                                 *
*  DEFINE ESCON CHPIDS                                           *
*                                                                 *
*****
      CHPID PATH=(70),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(71),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(72),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(73),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(74),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(75),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(FC),PARTITION=(HOST1),SWITCH=50,TYPE=FCV
      CHPID PATH=(FD),PARTITION=(HOST1),SWITCH=50,TYPE=FCV
*****
*  DEFINE 2107-9A2 LOGICAL CONTROL UNIT 0                         *
*                                                                 *
*****
      CNTLUNIT CUNUMBR=7000,PATH=(70,71,72,73),UNITADD=((00,256)), *
      LINK=(24,2D,34,3D),CUADD=0,UNIT=2107
*****
*  DEFINE 3390-9 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 0 *
*      16 BASE ADDRESS, 3 ALIASES PER BASE                       *
*****
      IODEVICE ADDRESS=(9000,016),CUNUMBR=(7000),STADET=Y,UNIT=3390B
      IODEVICE ADDRESS=(90D0,048),CUNUMBR=(7000),STADET=Y,UNIT=3390A
*****
```

```

* DEFINE 2107-9A2 LOGICAL CONTROL UNIT 1 *
*
*****
CNTLUNIT CUNUMBR=7001,PATH=(70,71,72,73),UNITADD=((00,128)), *
LINK=(24,2D,34,3D),CUADD=1,UNIT=2107
*****
* DEFINE 3390-9 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 1 *
* 16 BASE ADDRESS, 3 ALIASES PER BASE *
*****
IODEVICE ADDRESS=(9100,016),CUNUMBR=(7001),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(9150,048),CUNUMBR=(7001),STADET=Y,UNIT=3390A
*****

```

```

*****
* DEFINE 2107-9A2 LOGICAL CONTROL UNIT 2 *
*
*****
CNTLUNIT CUNUMBR=7002,PATH=(70,71,72,73),UNITADD=((00,256)), *
LINK=(24,2D,34,3D),CUADD=2,UNIT=2107
*****
* DEFINE 3390-3 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 2 *
* 64 BASE ADDRESS, 3 ALIASES PER BASE *
*
*****
IODEVICE ADDRESS=(9200,064),CUNUMBR=(7002),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(9240,192),CUNUMBR=(7002),STADET=Y,UNIT=3390A
*****

```

```

* DEFINE 2107-9A2 LOGICAL CONTROL UNIT 3 *
*
*****
CNTLUNIT CUNUMBR=7003,PATH=(70,71,72,73),UNITADD=((00,256)), *
LINK=(24,2D,34,3D),CUADD=3,UNIT=2107
*****
* DEFINE 3390-3 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 3 *
* 64 BASE ADDRESS, 3 ALIASES PER BASE *
*
*****
IODEVICE ADDRESS=(9300,064),CUNUMBR=(7003),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(9340,192),CUNUMBR=(7003),STADET=Y,UNIT=3390A
*****

```

```

* DEFINE 2107-9A2 LOGICAL CONTROL UNIT 4 *
*
*****
CNTLUNIT CUNUMBR=7004,PATH=(74,75,FC,FD),UNITADD=((00,256)), *
LINK=(24,2D,34,3D),CUADD=4,UNIT=2107
*****
* DEFINE 3390-3 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 4 *
* 32 BASE ADDRESS, 7 ALIASES PER BASE *
*
*****
IODEVICE ADDRESS=(9400,032),CUNUMBR=(7004),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(9420,224),CUNUMBR=(7004),STADET=Y,UNIT=3390A
*****

```



```

*   DEFINE 2107-9A2 LOGICAL CONTROL UNIT 5                               *
*                                                                           *
*****
CNTLUNIT CUNUMBR=7005,PATH=(74,75,FC,FD),UNITADD=((00,256)),          *
LINK=(24,2D,34,3D),CUADD=5,UNIT=2107
*****
*   DEFINE 3390-3 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 5    *
*   32 BASE ADDRESS, 3 ALIASES PER BASE                                  *
*                                                                           *
*****
IODEVICE ADDRESS=(9500,032),CUNUMBR=(7005),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(95A0,096),CUNUMBR=(7005),STADET=Y,UNIT=3390A
*****

```

Guidelines for defining the logical control unit for the DS8000

You can use these guidelines to define the logical control unit for the DS8000.

You need the following information in order to define the logical control unit (LCU):

- Which devices (arrays) are 3380 format, and which are 3390 format
- Which devices have PAVs, and which do not
- How many devices you define to a given LCU

Note:

1. An LCU is equivalent to a logical subsystem (LSS). Like an LSS, an LCU can have a maximum of 256 logical devices or volumes.
2. Every device has a single unit address (in range 00 - FF).
3. An LCU can manage a maximum of 256 devices and identifies a device based on the unit address.

Update the I/O configuration with the appropriate devices, using the following guidelines:

- Three device types are possible for 3380: 3380, 3380B, and 3380A.
- Three device types are possible for 3390: 3390, 3390B, and 3390A.
- Any non-PAV 3380 is a 3380, regardless of model.
- Any non-PAV 3390 is a 3390, regardless of model.
- Any 3380 PAV base is a 3380B.
- Any 3390 PAV base is a 3390B.
- Any alias of a 3380B, regardless of model, is a 3380A.
- Any alias of a 3390B, regardless of model, is a 3390A.
- Only define the actual base and alias devices to IOCP or HCD to avoid unnecessary definition of processor subchannels.
- Only define the unit address range to IOCP or HCD as necessary for the installed devices in an LCU (64, 128, or 256). Defining the unit address range enables you to put more LCUs and devices on a channel path.

IBM recommends that you document the configuration you define for the DS8000. With the exception of the 3380 and 3390 models, the DS8000 configuration and the IOCP input must match.

Note: The CU device is a 2107 for any DS8000 device type. In HCD you can use 3390 as a control-unit type for 3380 and 3390 devices. However, you must use 2107 as the control-unit type for 3380B, 3380A, 3390B, and 3390A devices.

Reconfiguring an LCU with larger capacity 3390 volumes

Use the following process to reconfigure an LCU with larger capacity 3390 volumes.

To retain your data, you must offload the data from the address range of the ranks to be reformatted by performing the following steps:

1. Confirm the raid ranks that are associated with the address range that is being modified.
2. Vary the address range offline. This causes any associated parallel access volumes (PAVs) to become unbound.
3. Issue Devserv (DS QP,xxxx,nnnn) using the alias addresses to confirm aliases are unbound.
4. Undefine the raid rank to allow the configuration of larger capacity volumes.
5. Reassociate the PAVs.
6. Vary the base address range online.
7. Issue Devserv (DS QP,xxxx,nnnn) using the base addresses to confirm that the aliases are bound.
8. Restore the data that you offloaded.

Increasing the number of 3390s formatted within an LSS

Use this process to increase the number of 3390s that are formatted within an LSS.

To increase the number of 3390s that are formatted within an LSS, perform the following:

1. Define an input/output definition file (IODF) using HCD with the desired device addresses.
2. Remove the parallel access volumes (PAVs) using the DS8000 CLI.
3. Issue Devserv (DS QP,xxxx,nnnn) to the alias addresses to initiate an I/O operation.

Note: Messages "IOS017I ALIAS DEVICE adev IS UNBOUND" and "IOS001E xxxx, INOPERATIVE PATH" are issued from all the alias addresses. These messages are normal and are not cause for concern.

4. Activate a new IODF by either performing an IPL or by dynamically activating the new IODF.

Attention: Ensure that there is sufficient memory in the hardware storage area to store the new configuration.

5. Add new 3390 volumes to the LCU.
6. Assign new PAVs.

Copy Services

This topic provides overviews of the different functions of Copy Services, including the benefits of using Copy Services for disaster recovery.

Copy Services overview

Copy Services is a collection of functions that provide disaster recovery, data migration, and data duplication functions. Copy Services runs on the IBM TotalStorage DS8000 storage unit and supports open systems and zSeries environments.

Many design characteristics of the DS8000 and the data copying and mirroring capabilities of the Copy Services features contribute to the protection of your data, 24 hours a day and seven days a week. The following information provides a brief description of each of these licensed features.

You can manage Copy Services functions through the command-line interface called the IBM TotalStorage DS CLI and the Web-based interface called the IBM TotalStorage DS Storage Manager. The DS Storage Manager allows you to set up and manage the following data-copy features from anywhere that network access is available.

Point-in-time copy

The point-in-time copy feature, which includes FlashCopy, enables you to create full volume copies of data using source and target volumes that span logical subsystems within a single storage unit. After the FlashCopy function completes, you can immediately access both the source and target copies.

Remote Mirror and Copy

The remote mirror and copy (previously known as Peer-to-Peer Remote Copy or PPRC) feature copies data between volumes on two or more storage units. When your host system performs I/O operations, the source volume is copied (mirrored) to the target volume automatically. After you create a remote mirror and copy relationship between a source volume and a target volume, the target volume continues to be updated with changes from the source volume until you remove the relationship between the volumes.

The following functions of remote mirror and copy are available:

- IBM TotalStorage Metro Mirror (previously known as Synchronous PPRC)
- IBM TotalStorage Global Copy (previously known as PPRC Extended Distance)
- IBM TotalStorage Global Mirror (previously known as Asynchronous PPRC)

Remote mirror for z/OS

z/OS Global Mirror (previously known as Extended Remote Copy or XRC) provides a long-distance remote copy solution across two sites for open systems and z/OS data using asynchronous technology. It is based on the existing Copy Services functions called Global Copy and FlashCopy and additional new functions that form consistency groups across one or more storage units to manage and control Global Mirror sessions.

Consider the following Copy Services information:

- The licenses for Copy Services features (remote mirror and copy and point-in-time copy) must be activated in the DS Storage Manager before you can begin using the features.
- A Copy Services domain, configured as part of the IBM TotalStorage Enterprise Storage Server® (ESS) Models 750 and 800, is supported on the DS8000 storage unit.

IBM TotalStorage Resiliency Family

IBM TotalStorage Resiliency Family is a set of products and features that are designed to help you implement storage solutions that keep your business running 24 hours a day, 7 days a week.

The IBM TotalStorage Resiliency Family is a set of hardware and software features, products, and services that are available on the IBM TotalStorage DS8000 and IBM TotalStorage Enterprise Storage Server Models 750 and 800. The DS8000 has a number of advanced Copy Services functions that are part of the IBM TotalStorage Resiliency family. Copy Services is a collection of functions that provides disaster recovery, data migration, and data duplication functions. Copy Services runs on the DS8000 storage unit and supports open systems and zSeries environments.

The following are key features of the IBM TotalStorage Resiliency family:

- Point-in-time copy, which includes FlashCopy
- Remote mirror and copy, which includes the following functions:
 - IBM TotalStorage Metro Mirror
 - IBM TotalStorage Global Copy
 - IBM TotalStorage Global Mirror
- Remote mirror and copy for z/OS, which includes z/OS Global Mirror and z/OS Metro/Global Mirror, and is available only for zSeries environments

Disaster recovery using Copy Services

The prime purpose for backing up, copying, and mirroring data is in preparation for a possible disaster.

Having a disaster recovery plan can ensure that critical data is recoverable at the time of a disaster. Because most disasters are unplanned, your disaster recovery plan must provide a way that allows you to recover your applications quickly, and more importantly, to access your data. This means that you must be able to restart your critical applications at a backup site.

Disasters often occur over a period of time. This type of unplanned disaster, called a rolling disaster, is the most common. Other unplanned outages include fatal hardware failures or data center damage that is caused by natural occurrences, such as earthquakes or hurricanes. Planned outages include system maintenance operations, and although they are intended, planned outages can be just as effective as unplanned outages at removing service from users.

Most users decide to use a combination of remote mirror and copy and point-in-time copy features to form a comprehensive enterprise solution for disaster recovery.

Disaster recovery solutions include, but are not limited to the following functions:

Global Mirror

Global Mirror provides recoverable copies of your production data at a long distance remote site by continually forming consistent sets of data with no significant performance impact. This allows a quick restart at your remote site in case of a disaster at your production site.

Failover and failback operations

A failover operation is the process of switching production to a remote site. A failback operation is the process of returning production to its original location after a disaster. These operations help reduce the time that is required to synchronize remote mirror and copy volumes after switching between sites during planned and unplanned outages.

Copy Services functions across a 2105 and 2107

Copy Services functions that are performed using either the DS Storage Manager or DS CLI are designed to interact with both the DS8000 (2107) and the IBM TotalStorage 2105 Enterprise Storage Servers (ESS) Models 800 and ESS 750.

The 2107 is the next generation of the Enterprise Storage Server (ESS). Most Copy Services functions that are available on the 2105 are also available in the 2107 and in open systems and zSeries environments.

On the 2107, a storage unit partitioning consists of two logical partitions (LPARs), one LPAR on each processor complex and a set of allocated resources such as processors, memory, I/O adapters, and storage devices in one physical machine. Each LPAR appears to the software running within it as an independent hardware environment. For this reason, the mirroring solutions, which are compatible between the 2105 and the 2170 series, allow you to set up a disaster recovery solution with the 2107 at your local (or primary) site and with the 2105 at your secondary (or recovery) site, where the same performance is not required.

To perform Copy Services functions between machine types 2105 and 2107, you must configure a Copy Services domain on the DS Storage Manager or DS CLI. To connect to the 2105 Copy Services domain on the ESS, all interfaces that you use require an authenticated login procedure to access Copy Services functions across the storage complex. The authentication is performed by using a user name and password that was created with the ESS Specialist. Therefore, the existing user name and password that was created with the ESS Specialist for the 2105 Copy Services domain for which you will be working must match the user name and password on the management console that is connected to the 2107. Otherwise, you must add them using either the DS Storage Manager or DS CLI as part of the procedure for adding a 2105 Copy Services domain to the storage complex.

Consider the following guidelines:

- To manage Copy Services across the 2105 and 2107, you must install licensed internal code version 2.4.2 or later on the 2105.
- The 2107 support remote mirror and copy (formerly PPRC) operations with an ESCON link. Therefore, if you want to configure a remote mirror and copy relationship between the 2107 and 2105, you must use a FCP link.

Adding a Copy Services 2105 domain to a storage complex using a DS Storage Manager

This process describes the information that you need to set up a Copy Services 2105 domain using the DS Storage Manager.

Copy Services functions are available through the DS Storage Manager that is installed on a management console. The management console is the focal point for configuration and Copy Services functions.

Configuring a Copy Services domain on the DS Storage Manager allows you to perform Copy Services functions between machine type 2105 and machine type 2107. The mirroring solutions are compatible between the 2105 (ESS Models 750 and 800) and 2107 .

For this example, assume that you want to configure a 2105 Copy Services domain using the DS Storage Manager. The 2107 must authenticate with the 2105 before you can perform Copy Services functions. Because of the single login process used by the 2107 to access the 2105, the user names and passwords must be the same

on the 2107 and 2105. This allows the user to access the 2105 server from the DS Storage Manager, without having to authenticate on the 2105.

As the administrator, perform the following steps to configure the 2105 Copy Services domain and add it to the list of storage complexes.

1. **Determine the address of the management console.** You need to know the IP address or fully qualified host name of the management console for one of the servers in the Copy Services domain. If you know the address, you can ping the name of the machine. If you do not, see your network administrator. You will be providing this address in Step 3.
2. **Identify the user name and password of that management console.** You need this information to access information from the storage complex. For example, depending on which client application the DS Storage Manager is running (local host or management console), you need the existing user name and password that was created with the ESS Specialist for the 2105 Copy Services domain. The user name and password must match on the management console and 2105 Copy Services domain in order to connect. If they do not match, you must create the user account (in Step 3) that matches the user account on the 2105.

Note: The Storage Complex drop-down list will include *local host* as an entry if you installed a DS Storage Manager on your own workstation.

3. **Add a connection for a 2105 Copy Services domain.**
 - a. To add a connection for the 2105 Copy Services domain, expand the **Manage Hardware** section, click **Storage Complexes**, and click **Add 2105 Copy Services Domain** and click **Go**.
 - b. Add the IP address of the 2105 Copy Services IP address and click **OK**. This is the IP address that you identified in Step 1. If the user name and password on the management console and 2105 Copy Services domain match, the status is shown as connected. If they do not match, continue to the following step.
 - c. Create a user account that matches the user account from the 2105 Copy Services domain. To do this, expand the **Monitor System** section, click **User Administration**, and click **Add user**.
 - d. From the **Add user** page, enter the user name and a temporary password. You are asked which access or group assignment that you want to assign the user. For this example, select the Copy Services group to allow the user to manage Copy Services relationships.

Copy Services operator

Performs Copy Services functions

- e. Logout from the DS Storage Manager
- f. Log back into the DS Storage Manager using the user name that matches the user account from the 2105. Then, change the password to match the password on the 2105 Copy Services domain. You can now create Copy Services relationships between the 2105 and 2107 .

Creating a Metro Mirror volume pair between a 2107 and a 2105

This process describes how to create a Metro Mirror volume pair using volumes from a 2107 and a 2105.

Before you begin, ensure that you meet the following requirements:

- The license for the remote mirror and copy feature must be activated.

- To create a Metro Mirror volume pair between Model types 2107 and a 2105, you must have added the **Add 2105 Copy Services Domain** option from the Section Action drop-down list when you configured your storage complex environment.
- Ensure that paths are set up between the source and the target LSSs for the Metro Mirror volume pairs. The paths between the 2105 and the 2107 must be configured using Fibre Channel Protocol (FCP) ports.
- The storage type of the source and target volumes on the 2107 and 2105 domain must have the same type. That is, if the source volumes are fixed block volumes, the target volumes must also be fixed block volumes.
- The size of the volumes in the source LSS must be less than or equal to those of the target LSS.

You can create Metro Mirror relationships using source and target volumes from the following machine types:

- A 2107 and a 2107
- A 1750 and a 1750
- A 2107 and a 1750
- A 2105 and a 1750
- A 2105 and a 2107

Note: If the source is a Copy Services 2105 domain, the Metro Mirror task is performed on the source domain. However, if you perform a "Suspend at target" action, the suspension occurs at the target domain.

To create a Metro Mirror pair between a 2107 and a 2105, perform the following steps. For this task, the source domain is a 2105 Model 800 or 750 and the target is a 2107.

1. In the navigation, under Copy Services, click Metro Mirror.
2. Select the **2105 Copy Services domain** from the list of storage complexes. This option is enabled if you included a Copy Services domain for the 2105 when you selected Storage Complex from the Section Action drop-down list under **Manage Hardware**. All storage complexes that you added when you configured your environment are included in the drop-down list.
3. Select a storage unit that you will be working with from the **Storage unit** drop-down menu.
4. Select a storage image from the **Storage image** drop-down menu.
5. Select a resource type from the **Resource type** drop-down menu.
6. Select which resource to view from the **Specify (LSS, Volume Group, Host attachment, Storage type)** drop-down menu.
7. Click on **Select Action** drop-down menu and select **Create...** and click **Go**.
8. On the Volume pairing method page, specify one of the following methods by which to select volume pairs and then click **Next**.
 - a. If you click the **Automated volume pair assignment**, the first selected source volume is paired automatically with the first selected target volume.
 - b. If you click **Manual volume pair assignment**, you must select each specific target volume for each selected source volume.
9. On the Select source volumes page, select the volumes that you want to include for the Metro Mirror relationships and click **Next**. If the resource type is an LSS, specify which from LSS to select source volumes. Optionally, you can click the **Create paths** button to create a new path. You are linked to the page where you can create a path.

10. On the Select target volumes page, select target volumes by changing the storage complex to the name that is associated with the 2107 domain, and click **Next**. Consider the following guidelines:
 - a. If you selected the automated pairing method, each source volume that you selected is paired with a target volume.
 - b. If you selected the manual pairing method, you can select target volumes from different LSSs on the target storage unit. This page repeats for each of the source volumes that you selected.
11. On the Select copy options page, select the type of relationship and the copy options for the relationship and click **Next**. The **Enable critical volume mode** and **Create relationship even if target is online to host** options are available for count-key-data volumes only.
 - a. Select the **Perform initial copy** option the first time a that volume pair is created because it is needed to ensure that the entire source volume is copied to the target volume to guarantee that the source and target volumes contain the same data.
 - b. Optionally, you can select the **Suspend Metro Mirror relationship after initial copy** option because it only applies to Metro Mirror volume pairs. This means that after data is completely copied to the target volume of the volume pair, the target volume goes into the suspended state.
12. On the Verification page, review the attributes and values to verify that they are correct.
13. If the attributes and values are not correct, click **Back** to return and specify the correct values. Otherwise, click **Finish** to complete the process.

FlashCopy

The FlashCopy function enables you to make point-in-time, full volume copies of data, with the copies immediately available for read or write access. You can use the copy with standard backup tools that are available in your environment to create backup copies on tape.

FlashCopy creates a copy of a source volume on the target volume. This copy is called a point-in-time copy. When you initiate a FlashCopy operation, a FlashCopy relationship is created between a source volume and target volume. A FlashCopy relationship is a "mapping" of the FlashCopy source volume and a FlashCopy target volume. This mapping allows a point-in-time copy of that source volume to be copied to the associated target volume. The FlashCopy relationship exists between this volume pair from the time that you initiate a FlashCopy operation until the storage unit copies all data from the source volume to the target volume or you delete the FlashCopy relationship, if it is a persistent FlashCopy.

The benefits of FlashCopy are that the point-in-time copy is immediately available for use for backups and that the source volume is immediately released so that applications can be restarted, with minimal application downtime. The target volume is available for read and write processing so it can be used for testing or backup purposes.

When the data is physically copied, a background process copies tracks from the source volume to the target volume. The amount of time that it takes to complete the background copy depends on the following criteria:

- The amount of data being copied
- The number of background copy processes that are occurring

- The other activities that are occurring on the DS8000

FlashCopy supports the following copy options:

Consistency groups

Creates a consistent point-in-time copy of multiple volumes, with negligible host impact. You can enable FlashCopy consistency groups from the command-line interface (DS CLI).

Change recording

Activates the change recording function on the volume pair that is participating in a FlashCopy relationship. This enables a subsequent refresh to the target volume.

Establish FlashCopy on existing Metro Mirror source

Allows you to establish a FlashCopy relationship where the target volume is also the source of an existing remote mirror and copy source volume. This enables you to create full or incremental point-in-time copies at a local site and then use remote mirroring commands to copy the data to the remote site.

Fast reverse

Reverses the FlashCopy relationship without waiting for the finish of the background copy of the previous FlashCopy. This option applies to the Global Mirror mode.

Inhibit writes to target

Ensures that write operations are inhibited on the target volume until a refresh FlashCopy operation is complete.

Multiple Relationship FlashCopy

Allows a source volume to have multiple target volumes at the same time.

Persistent FlashCopy

Allows the FlashCopy relationship to remain even after the FlashCopy operation completes. You must explicitly delete the relationship.

Refresh target volume

Provides the ability to “refresh” a FlashCopy relationship, without recopying all tracks from the source volume to the target volume.

Reverse restore

Reverses the FlashCopy relationship and copies data from the target volume to the source volume.

Reverse FlashCopy relationships

The direction of a FlashCopy relationship can be reversed, in which case the volume that was previously defined as the target becomes the source for the volume that was previously defined as the source (and is now the target). The data that has changed is copied to the volume previously defined as the source.

You can reverse a FlashCopy relationship if you want to restore a source volume (Volume A) to a point-in-time before you performed a FlashCopy operation. In effect, you are reversing the FlashCopy operation so that it appears as though no FlashCopy operation ever happened. Keep in mind that the background copy process of a FlashCopy operation must complete before you can reverse volume A as the source and volume B as the target.

There might be certain circumstances when you might need to reverse an original FlashCopy relationship. For example, suppose you create a FlashCopy relationship

between source volume A and target volume B. Data loss occurs on source volume A. To keep applications running, you can reverse the FlashCopy relationship so that volume B is copied to volume A.

Note: A fast reverse option that applies to a Global Mirror operation allows a FlashCopy relationship to be reversed without waiting for the background copy of a previous FlashCopy relationship to finish. A Global Mirror operation is based on existing Global Copy and FlashCopy operations at the target site.

Figure 8 illustrates how a reverse restore operation works:

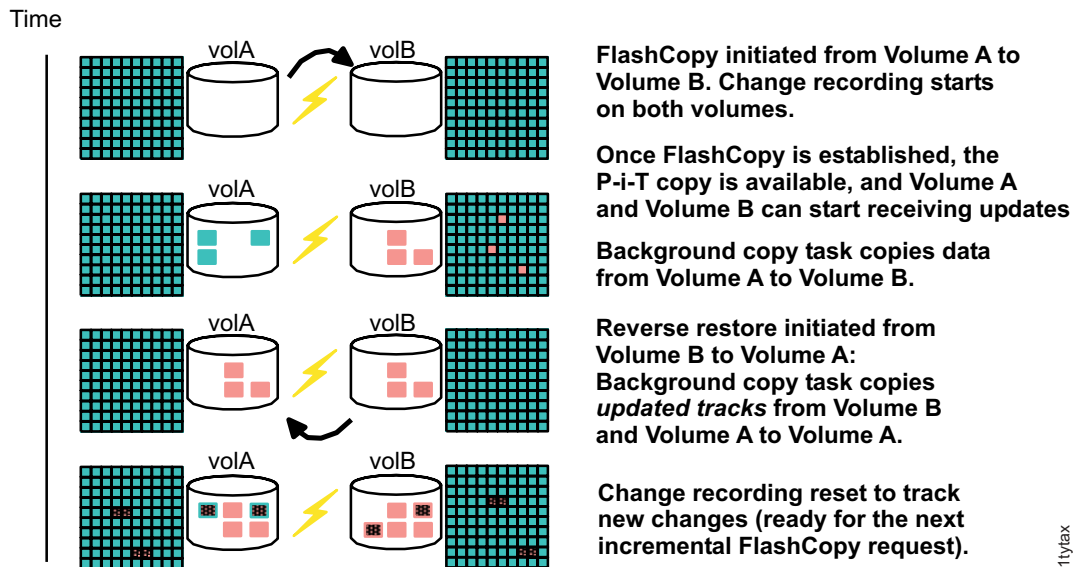


Figure 8. Refreshing target volume — reverse restore

FlashCopy to Metro Mirror source volumes

You can use an existing Metro Mirror source volume as a FlashCopy target volume. This process allows you to create a point-in-time copy and then make a copy of that data at a remote site.

Previously, you could create a FlashCopy relationship between two volumes and then create a Metro Mirror volume pair by using the FlashCopy target volume as the Metro Mirror source volume. However, you could not use an existing Metro Mirror source volume as a FlashCopy target volume. This is no longer a restriction. You can create FlashCopy target volumes that serve as production volumes and mirror them to a remote site using Metro Mirror.

You can create a full volume or incremental point-in-time copy to an existing Metro Mirror source volume. The first time that you create a FlashCopy operation, the entire Metro Mirror source volume has to be copied to the remote site, similar to a new pair being established. Afterwards, provided that you created the FlashCopy volume pair with the persistent and change recording options, you can perform a refresh (incremental) of the FlashCopy relationship. The refresh copies only changes and sends them to the remote site, requiring less time for the remote mirror and copy volumes to be synchronized again.

When you create a FlashCopy operation and the logical copy finishes, the corresponding Metro Mirror volume pairs turn to duplex pending status as long as

the FlashCopy data is being transferred to the Metro Mirror remote site. The time it takes to complete the copy of the FlashCopy data, until all remote mirror and copy volumes are synchronous again, depends on the amount of data to be transferred and the available bandwidth to the remote site. Be aware that, during this time, the remote site will be inconsistent if a disaster occurs. Therefore, you should carefully decide if you can tolerate not mirroring FlashCopy target volumes temporarily at a remote site.

Multiple relationship FlashCopy

Multiple relationship FlashCopy allows a source volume to have multiple targets simultaneously.

When a FlashCopy operation copies data from a source volume to a target volume, that source volume can be involved in more than one FlashCopy relationship at the same time (known as a multiple relationship FlashCopy). If a track on a volume is not a target track of an existing FlashCopy relationship, it can become the target in a new FlashCopy relationship. Up to twelve different target FlashCopy relationships can exist for a single source volume.

If you require multiple copies of the same data, this function allows you to copy a single source volume multiple (up to twelve) times to different target volumes as soon as a FlashCopy volume is established. For example, suppose you use FlashCopy to copy volume A to volume B. As soon as that FlashCopy relationship is established, you can copy volume A to volume C. Once that relationship is established, you can copy volume A to volume D, and so on. Multiple sessions within a single volume are also possible.

Persistent FlashCopy

A persistent FlashCopy relationship is one that does not end when the background physical copy ends. The relationship persists until explicitly withdrawn.

Persistent FlashCopy overrides normal FlashCopy processing and allows a FlashCopy relationship to stay active until explicitly withdrawn. Generally, a FlashCopy relationship is automatically withdrawn when the background physical copy ends.

Designating that a FlashCopy relationship be persistent can help protect against inadvertent updates of recently created target volumes. For example, if a source volume is regularly copied to alternating target volumes (thereby ensuring that a complete copy of the source volume is always available), the persistent relationship will identify the target volume for the most recently completed FlashCopy.

Consider the following situation:

- A company wants to maintain more than one point-in-time copy version of a set of its production volumes.
- The company creates a FlashCopy of its production volumes every eight hours, each time using a different set of target volumes.
- The production volume is Volume A.
- The copies are Volume B, Volume C, and Volume D.
- The volumes that are copied contain a different point-in-time copy version.

The following represents the procedure that FlashCopy operations follow in this scenario:

1. FlashCopy from Volume A to Volume B
2. FlashCopy from Volume A to Volume C
3. FlashCopy from Volume A to Volume D

This FlashCopy procedure presents a possible exposure for mistakes. For instance, if the FlashCopy operation (step 2 of the FlashCopy procedure) is mistakenly performed from Volume A to Volume B rather than from Volume A to Volume C, the result could be data that overwrites the original point-in-time copy in volume B.

The use of the persistent FlashCopy function requires that the FlashCopy relationship is explicitly withdrawn before another FlashCopy operation can be performed. While this restriction does not guarantee that a point-in-time copy cannot be overwritten by mistake, it does provide an extra measure of security.

Note: Persistent FlashCopy is always required for refreshing a FlashCopy target volume.

Re-synchronizing FlashCopy volume pairs

Re-synchronizing FlashCopy volumes provides the ability to update an initial point-in-time copy of a source volume.

You can resync your FlashCopy volumes without having to recopy your entire volume. This process is only possible with a persistent relationship, whereby the storage unit continually tracks updates to the source and target volumes. With persistent relationships, the relationship between the source and target volumes is maintained after the background copy has completed.

This function reduces the background copy completion time because only a subset of data is copied when either the source or target volume has changed. The target volume becomes current with the newly established point-in-time source copy.

To enable this function of FlashCopy, the storage unit uses a feature called change recording to monitor write operations and record changes that were made to volumes participating in FlashCopy relationships since the initial (or last) incremental FlashCopy was performed. After the initial copy of the source volume to the target volume, the FlashCopy relationship between the volumes is not automatically deleted. Rather, it remains to allow the volumes to be re-synchronized. The next time that you want to update the information on your target volume, you can resync your target volume again.

Note: If you perform a resync operation, it is sent to any source and volume pair. All source and target volume pairs will be re-synchronized, not just the one pair that you selected.

Creating a Remote FlashCopy operation (Inband FlashCopy on the ESS 2105)

Remote FlashCopy operations (Inband FlashCopy on the ESS 2105) allow you to create a FlashCopy relationship at a target (remote) site using remote FlashCopy commands, without having an network connection to the target site.

Remote FlashCopy commands are issued to a source volume of a remote mirror and copy volume pair on a source (local) storage unit and sent across paths (acting

as a conduit) to a target storage unit to establish a FlashCopy relationship at the target site. This eliminates the need for a network connection to the target site solely for the management of FlashCopy.

Limitation: Remote FlashCopy commands are used to establish a FlashCopy relationship at the target (remote) site when a connection to target site is lost. The Remote FlashCopy operation is not supported through the DS Storage Manager because connections to both the source and target sites are required. If the connection to the target site is lost, the DS Storage Manager cannot connect to the target site.

The following example illustrates the required steps for a Remote FlashCopy operation. You can perform all steps from the DS CLI. If you used the DS Storage Manager for Steps 1 and 2, you must perform Step 3 from the DS CLI.

1. **Create paths between the source LSS and the target LSS.** For example, IBM.2107-1600661/00 and IBM.2107-1306875/00
2. **Create Metro Mirror volume pairs from the source LSS to the target LSS** for example, volume A (IBM.2107-1300861/0001) from LSSA and volume B (IBM.2107-1300871/0001) from LSSB.
3. **Enable a Remote FlashCopy operation at the target site using volume B as the source volume and volume C as the target volume** Assume that the target site connection is lost. You can create the FlashCopy relationship from volume B to volume C (both from the target site). However, you cannot use the DS Storage Manager for this step because connections to the target site are lost.

Assume that you performed Step 1 and Step 2 from the DS Storage Manager (connections to both storage units at the source and target sites were available) and the Metro Mirror relationship between the volume pair still exists. To create the Remote FlashCopy operation, you must perform Step 3 from the DS CLI using the following command as an example:

```
dsccli>mkremoteflash -dev IBM.2107-75FA120 -conduit IBM.2107-75FA150/100100:0200
```

where:

dev Specifies the storage image ID, which includes manufacturer, type, and serial number.

conduit LSS_ID

(Required) Identifies the source LSS of an existing remote mirror and copy relationship that is used as a conduit for communicating with the target storage image. The source volume IDs that are specified in *source_volume_ID:target_volume_ID* must be the target volumes in a remote mirror and copy relationship in which one of the conduit LSS volumes acts as a source volume. You can specify a fully qualified LSS ID, which includes the storage image ID.

source_volume_ID:target_volume_ID

(Required) Establishes a remote FlashCopy relationship for the source and target volume pairs with the specified IDs. You can specify fully qualified volume IDs, which includes storage image IDs, or a shortened version without storage image IDs if the -dev parameter is specified. Separate the IDs of the FlashCopy relationships with spaces.

Path connections

Before you can create source and target remote mirror and copy copy pair relationships, you must establish logical paths between a logical subsystem (LSS) in a source storage unit and an LSS in a target storage unit.

When you implement a remote mirror and copy configuration, the definition of paths that remote mirror and copy uses is an important consideration. To establish paths, you must have Fibre Channel Protocol (FCP) adapters on both the source and target storage unit for the peer-to-peer links. Paths are required for communicating between volume pairs and for copying data from a source volume to a target volume.

You can establish FCP paths to connect to a source LSS and a target LSS that contain count-key-data (CKD) or fixed block (FB) volumes. The requirement is that both the source LSS and target LSS must contain the same type of volumes.

You can define from one to eight FCP paths, depending on the capability of the LSS, from a single source site LSS to a specific target site LSS.

When you establish remote mirror and copy paths over FCP links, the source storage unit uses a worldwide port name (WWPN) to communicate with the target storage unit. The worldwide node name (WWNN) and system adapters identification (SAID) of the source and destination adapters are used to generate a WWPN. A WWPN consists of exactly 16 hexadecimal digits. The hexadecimal characters are 0–9, a–f, and A–F. The values for the WWNN and the WWPN must be unique.

With FCP, a single fibre channel link between two ports enables bi-directional remote mirror and copy operations. This means that it is possible to send and receive data in both directions simultaneously. For example, one link can have a path that is established in one direction and then have another path that is established in another direction at the same time, on the same physical path.

While a FCP interface can accommodate simultaneous data transfers, it does have limited bandwidth. To distribute workload evenly to all available paths, the storage unit monitors the overall workload on each port and selects paths that are determined by the size of the data transfer, the available bandwidth available on each FCP port, and the number of data transfers, currently processing on each port. Selecting paths in this manner ensures good response time and overall system throughput.

Remote Mirror and Copy

Remote mirror and copy is a storage server feature that constantly updates a target copy of a volume so that it matches changes that were made to its associated source volume.

The remote mirror and copy feature is a hardware solution that enables the mirroring of data from one site (the local site) and its associated volumes (source volumes) to a second storage unit at another site (the recovery or remote site) and its volumes (target volumes).

Note: Remote mirror and copy was referred to as Peer-to-Peer Remote Copy (PPRC) in earlier documentation for the IBM TotalStorage Enterprise Storage Server. Currently, the remote mirror and copy feature name is used generically for all methods of synchronous and asynchronous methods. In

addition, in a remote mirror and copy environment, a site where production applications run might be referred to as the production site, local site, or Site A. A site that has mirrored data of your local site might be referred to as a recovery site, backup site, remote site, or Site B.

Remote mirror and copy differs from FlashCopy in two essential ways. First, the source and target volumes can be on the same storage server or on separate storage servers that are located at some distance from each other. Second, and more significantly, remote mirror and copy does not capture the state of the source volume at some point in time, but rather reflects all changes that were made on the source volume to the target volume.

The following modes of operation are available:

IBM TotalStorage Global Copy

Global Copy copies data nonsynchronously and over longer distance than is possible with Metro Mirror. In Global Copy mode, the source volume sends a periodic, incremental copy of updated tracks to the target volume instead of a constant stream of updates. Because the original order of updates is not strictly maintained, this is considered a "fuzzy" copy. All updates that an attached host performs on the source volume are also performed on the target volume, but at a later time.

IBM TotalStorage Global Mirror

Global Mirror copying provides a two-site extended distance remote copy option for disaster recovery. This solution is based on existing Global Copy and FlashCopy. With Global Mirror, the data that the host writes to the storage unit at the local site is asynchronously shadowed to the storage unit at the remote site. A consistent copy of the data is then automatically maintained on the storage unit at the remote site.

IBM TotalStorage Metro Mirror

Metro Mirror synchronously mirrors updates that are done to source volumes. All updates that any attached host performs on the source volume are also performed on the target volume before operation completion is indicated. Metro Mirror copying supports a maximum distance of 300 km (186 mi).

Global Copy

Global Copy functions offer a nonsynchronous long-distance copy option whereby write operations to a storage unit at your production site are considered complete before they are transmitted to a storage unit at your recovery site.

Global Copy is a nonsynchronous mirroring function and is an alternative mirroring approach to Metro Mirror operations. Host updates to the source volume are not delayed by waiting for the update to be confirmed by a storage unit at your recovery site. The source volume sends a periodic, incremental copy of updated tracks to the target volume instead of a constant stream of updates. There is no guarantee that dependent write operations are transferred in the same sequence that they have been applied to the source volume. This nonsynchronous operation results in a "fuzzy copy" at the recovery site; however, through operational procedures, you can create a point-in-time consistent copy at your recovery site that is suitable for data migration, backup, and disaster recovery purposes.

To ensure that a consistent copy of the data is created, you can periodically switch from Global Copy to Metro Mirror mode. Then, either stop application I/O or freeze

write applications to source volumes and wait for all pending updates to be copied to the recovery site. At this point, you can create a FlashCopy operation at your recovery site to obtain a consistent data.

The Global Copy function can operate at very long distances—well beyond the 300 km distance that is supported for Metro Mirror—and with minimal impact to applications, with the distance that is limited only by the network and the channel extended technology.

During a disaster, data can be restored only to the last known consistent increment that was created. This means that data that is written to the production site but is waiting to be transferred to the recovery site is lost whenever the two storage units can no longer communicate. Be aware that the use of Global Copy functions do not guarantee against data loss. Global Mirror functions, on the other hand, provide recoverable copies of production data at a long distance recovery site by continually forming consistent sets of data with no significant performance impact. This allows a quick restart at your recovery site in case of a disaster at the production site.

The following describes the Global Copy write sequence:

1. During a Global Copy operation, the storage unit at your production site captures information about updates to the source and periodically sends those updates to the target volume at the recovery site.
2. After the initial copy of tracks, the storage unit series periodically starts a synchronization cycle where all updated tracks, in ascending order from the lowest numbered track, is copied from the source volume to the target volume. The storage unit updates the target tracks with the current information for that track, regardless of the number of updates between the time that it was last copied, the current time, and the order in which the updates occurred.
3. When this process completes, the cycle is repeated. There is little response time degradation on application write operations in extended distance mode.
4. Write updates to the source volume receive an immediate completion because the synchronization cycle is independent of the updates to the source volume.

Global Mirror

Global Mirror provides a long-distance remote copy solution across two sites for open systems or z/OS (or both) data using asynchronous technology.

Global Mirror provides a two-site extended distance remote copy option for disaster recovery. With Global Mirror, the data that the host writes to the storage unit at the production (local) site is asynchronously shadowed to the storage unit at the recovery site. The recovery site can be located well beyond the 300 km limit that is associated with Metro Mirror with minimal delays in response times. Separating volumes by larger distances provides greater protection against regional disasters. Only the capabilities of the network and channel extension technologies limit distance.

Global Mirror is based on existing Copy Services functions: Global Copy and FlashCopy. With Global Mirror, a consistent copy of the data is automatically maintained on the storage unit at the recovery site. This consistency is ensured if the order of dependent write operations to disks or disk groups is maintained. With Global Copy operations, data consistency is guaranteed at the recovery site by creating "freeze and run" tasks and then using FlashCopy to create point-in-time copies at the recovery site.

Global Mirror operations periodically invoke a point-in-time FlashCopy at the recovery site, at regular intervals, without disrupting the I/O to the source volume, thus giving a continuously updating, nearly up-to-date data backup. Then, by grouping many volumes into a session, which is managed by the master storage unit, you can copy multiple volumes to the recovery site simultaneously while maintaining point-in-time consistency across those volumes.

To better understand how Global Mirror works, you must be familiar with the following terms:

Master

The master storage unit controls the creation of consistency groups in the Global Mirror session. The master storage unit sends commands to subordinate storage units. A storage unit is allowed to be a master for only one Global Mirror session.

Subordinate

The subordinate storage unit receives commands from a master storage unit. The subordinate storage unit is identified when a Global Mirror session is started. The subordinate storage unit forms consistency groups and performs other Global Mirror processing. A subordinate storage unit can only be controlled by one storage unit master.

Session

A session is a collection of volumes across multiple storage units that are managed together to create consistent copies of data. The session is identified with an ID that is unique across the enterprise. The ID identifies volumes that will participate in the Global Mirror consistency group. A session is opened on each LSS in the enterprise that has or might have volumes that will participate in the Global Mirror consistency group that is associated with the specific session ID.

Control path

The control path is established from the master storage unit to the subordinate storage unit when more than one storage unit participates in the Global Mirror session. If there is only one storage unit that is involved, you do not have to create a control path. The master storage unit communicates directly with its subordinate storage unit.

Fast restore operations

The fast reverse option allows you to reverse a FlashCopy relationship without waiting for the background copy of a previous FlashCopy to finish.

In disaster recovery situations, when you use Global Mirror, you must use volumes with consistent data at the target site. Part of the setup requires that you enable the fast restore option, a copy option that is available for FlashCopy operations. (This option applies to a Global Mirror operation, which is based on existing Global Copy and FlashCopy operations at the target site.)

Use fast restore option to reverse a FlashCopy target volume and allow consistent data to be copied back to its associated source volume, without having to wait for the background copy from the original source to the original target to complete. You can then vary the volumes online and start your applications.

You can enable the fast reverse option by accessing Copy Services functions from the Web interface called the IBM TotalStorage DS Storage Manager and navigating to the FlashCopy page and then to the Select copy options page.

Global Mirror consistency groups

Global Mirror can control the formation of consistency groups for data consistency.

Global Mirror is based on a combination of Global Copy and FlashCopy functions. To support consistent data across storage units, Global Mirror uses a function called *sessions* for forming consistent copies. A consistency group is a collection of volumes across multiple storage units that are managed together when creating consistent copies of data. The formation of these consistency groups is coordinated by the master storage unit, which sends commands over remote mirror and copy links to its subordinate storage units.

With Global Mirror functions, consistency groups can be formed many times per hour, as opposed to once or twice per day. By combining many volumes into a session, which is managed by the Global Mirror master storage unit, multiple volumes can be copied to the recovery site simultaneously while maintaining point-in-time consistency across those volumes.

The following properties control how often consistency groups are formed. You can modify these values using the Web-based interface called the IBM TotalStorage DS Storage Manager. Access the Global Mirror page, and then Define Properties from the summary of actions menu.

Consistency group interval time

Indicates how long (in seconds) that the storage unit waits between the formation of consistency groups. If you set this value to zero (requires sufficient bandwidth), the consistency group forms continuously, which means that a consistency group starts forming as soon as the previous consistency group has been formed.

Maximum coordination interval

Indicates the maximum time (in milliseconds) that the master storage unit communicates with its subordinate storage units to form a consistent data point. The default is 50 milliseconds.

Because host write I/O is delayed while the consistent point is being set, performance might be affected by allowing too much time for this interval. If the time set for the maximum coordination interval expires before formation of the consistent data point is complete, the consistency group fails.

Maximum time writes are inhibited to the remote site

Indicates the maximum time (in seconds) that write operations are not allowed at the recovery site before the storage unit stops forming a current consistency group. If the drain time is maintained for an extended period of time, the number of write operations that are required to transfer the data to the recovery site can become large enough to increase the time to form a consistency group.

Note: As distance increases, there is an increasing lag in the time it takes for the data to be written to the recovery site. This lag is referred to as a *drain* time.

Loss of data can also be increased in the event of a disaster. If it is set to zero, it defaults to either four minutes or twice the consistency group interval value, depending on which value is larger. The first consistency group is formed regardless of the consistency group drain time. For the remainder of the consistency groups, if the specified time expires before the data is drained to the recovery site, consistency group formation stops.

After the consistency group formation has stopped five consecutive times, the timer is disabled and the next consistency group is formed regardless of the required time.

Session management

A Copy Services session is a collection of volumes across multiple storage units that are managed together to create consistent copies of data. Sessions are managed for Global Mirror relationships.

The Global Mirror mode supports data consistency across multiple volumes, logical subsystems (LSSs), and storage units by using a concept called *sessions*. For disaster recovery, it is required that the volumes in a session that are necessary for recovery be consistent up to the time of failure. If data on these volumes is copied out of sequence, then the recovery procedures that use these volumes become useless. Global Mirror uses special algorithms to provide data consistency.

Sessions are associated Global Mirror relationships and are defined with an identifier (session ID) that is unique across the enterprise. The ID identifies volumes that are related and that can participate in the Global Mirror consistency group.

You can select a set of volumes across LSSs on multiple storage units to participate in consistency groups. Logical volumes within a given LSS can be part of different sessions and different consistency groups. After you create a session, you can add logical volumes to or remove logical volumes from, a particular session. After a set of logical volumes become part of the session, a single session-based "command" can operate on all of the logical volumes in the session.

Metro Mirror

Metro Mirror is a function of the remote mirror and copy feature that constantly updates a secondary copy of a volume to match changes made to a source volume.

With Metro Mirror copying, the source and target volumes can be on the same storage unit or on separate storage units. You can locate the storage unit at another site some distance away. Synchronous mirroring means that each update to the source storage unit must also be updated in the target storage unit before another update can process. When Metro Mirror receives a host update to the source volume, it completes the corresponding update to the target volume. This guarantees data consistency by ensuring that a write operation that completes is received by the host application after the update has been committed to the target storage unit and acknowledged by both the source and target storage units. (Typically, the target volumes are on a different storage unit.) This results in near perfect data consistency but can result in lag time between transactions.

With Metro Mirror, consistency is guaranteed across all volumes on which an application does write operations as long as all volume pairs are in full duplex state. When error conditions affect some of the volume pairs (or different volume pairs at different time), this consistency might be lost. For example, if one of the target volumes cannot be updated because a path fails, the corresponding source volume normally goes into a suspended state, but still allows updates. However, these updates are no longer transferred to the target volume. Only the bitmap of changed tracks is created and maintained. So the consistency across volumes is lost, although the order of write operations is still guaranteed for the other target volumes.

Metro Mirror copying supports a maximum distance of 300 km (186 mi). Delays in response times for Metro Mirror are proportional to the distance between the volumes. However, 100% of the source data is available at the recovery when the copy operation ends.

The following procedure describes how data is written for a Remote Mirror and Copy operation. A copy to the target storage unit is synchronous with the source volume's I/O operation.

1. An application requests a write I/O to the source storage unit. The write I/O is written into cache and nonvolatile storage (NVS).
2. Metro Mirror sends the write I/O to the target storage unit cache and NVS.
3. The storage unit at the recovery site signals that the write operation has completed when the updated data is in its cache and NVS.
4. When the storage unit at the production site receives notification from the target storage unit that the write operation has completed, it returns the I/O completed status to your application.

Consistency groups

Consistency groups are used for controlling error situations and for maintaining data consistency at a recovery site.

You can create paths from a source LSS to a target LSS have the consistency group option enabled. All volume pairs in remote mirror and copy relationships from this source to target LSS (sharing the same paths) belong to the consistency group. Other volumes pairs are not affected.

With the consistency group option enabled, when an error occurs on any volume pairs or on the links that are associated with these LSS pairs, an alert is issued and I/O to all duplex remote mirror and copy volumes on LSS pairs will be queued either until a consistency group created operation is run or the consistency group timeout time expires, allowing external automation to use the consistency group created operation to create a dependent write consistent set of target volumes over any number of LSS and disk storage units.

External automation or command-line interface commands can issue a "freeze" to all LSS pairs that relate to the application, affecting the set of source volumes in the consistency group. The freeze operation causes the source volumes to suspend and their associated paths to end. Therefore, all associated updates at the recovery site are halted, resulting in a consistent data checkpoint. The automation program is not a part of the software provided with the storage unit; it must be supplied by the user. However, IBM has offerings to assist with this automation. For more information, contact your IBM storage representative.

A consistency grouping timeout function gives automation software the time to issue the automated requests to ensure consistency. The consistency grouping timeout is a value that specifies the length of time that the volume remains in long-busy (unavailable) following an error. You can change this value by accessing Copy Services functions from the Web interface (DS Storage Manager) and navigating to the Paths page and then to the LSS options page. The default value for this parameter is set to two minutes.

Note: The consistency group option is not required for Global Mirror or for Global Copy. Global Copy always keeps a inconsistent copy of the volumes at the recovery site. The volume pairs become consistent only when all updates

are transmitted and copied to the recovery site while the application write operations are quiesced. Global Mirror uses a function called a session to control the formation of consistency groups. (A session is a collection of volumes across multiple storage units that are managed together to create consistent copies of data.) The formation of these consistency groups is coordinated by the "master" storage unit, which sends commands over remote mirror and copy links to its "subordinate" storage units.

The following describes the behavior of these volume pairs when an error occurs that affects a volume in a consistency group:

- Without the consistency group option enabled, the storage unit causes the volume where the error is detected to enter a suspended state, but updates to that volume are still allowed.
- With the consistency group option enabled, the volume becomes suspended and additionally enters a long-busy state. Updates are not possible to volumes from the source LSS to target LSS that are sharing the same path.

A critical mode determines the behavior of the remote mirror and copy pairs or consistency groups after a failure in communication between the source and target storage units when all the paths between a pair are lost. You can enable this mode also by accessing Copy Services functions from the DS Storage Manager and by navigating to the Paths page and then to the LSS options page.

Failover and failback operations

The failover operation is the process of switching production to a backup facility (normally your recovery site). A failback operation is the process of returning production to its original location after a disaster or a scheduled maintenance period.

There are times, both planned and unplanned, when it is necessary to suspend disk mirroring and to make use of the secondary storage unit in your configuration. As a manual process, this can be complex. However, failover and failback recovery operations are available to simplify this process and reduce the risk of error and the time it takes to switch sites and restart I/O operations.

Failover is the process of temporarily switching production to a backup facility (normally your recovery site) following a scheduled maintenance period or a disaster at your production (or local) site. A failover operation is always followed by a failback operation, which is the process of returning production to its original location. These operations use remote mirror and copy functions to reduce the time that is required to synchronize volumes after switching sites during planned or unplanned outages.

The failover and failback operations allow change recording to be enabled on the target volumes without having to communicate between the target and source storage units. This method eliminates the need to perform a full volume copy from your recovery site to the production site, which can reduce the time that is required to resume operations at your production site.

In a typical remote mirror and copy environment, processing will temporarily failover to the storage unit at your recovery site if an outage occurs at the production site. Through use of failover operations, the state of a storage unit in your target configuration changes. As a result, the storage unit is recognized as the source storage unit in the pair. Because the failover process puts the volumes into a suspended state, changes are tracked within a bitmap. Assuming that change

recording is enabled, only change data is sent to the production site to synchronize the volumes, thereby reducing the time that is required to complete the failback operation.

When it is safe to return to your production site, assuming that no physical damage has occurred to the storage unit in the location, you can delete paths and create new ones from your production site to your recovery site. Then, you can create a failback recovery request to restore the storage unit as the production storage unit in the relationship.

The following considerations are for failover and failback operations:

- The failover operation does not reverse the direction of a remote mirror and copy pair. It changes a target device into a suspended source device, while leaving the source device in its current state.
- The failback recovery operation can be issued against any remote mirror and copy volume that is in a primary suspended state. The operation copies required data from the source volume to the target volume in order to resume mirroring. Failback operations are commonly used after a failover operation has been issued to restart mirroring either in the reverse direction (remote site to local site) or in the original direction (local site to remote site).

z/OS Copy Services

Copy Services functions are also supported in a z/OS environment. These functions allow you to create a point-in-time copy or to mirror data at your remote or recovery site as applications make changes to the source volumes at your production site.

z/OS Copy Services includes the following functions:

- z/OS Global Mirror
- z/OS Metro/Global Mirror

z/OS Global Mirror

z/OS Global Mirror (previously known as Extended Remote Copy or XRC) provides a long-distance remote copy solution across two sites for open systems and z/OS data using asynchronous technology.

DS8000 storage complexes support the z/OS Global Mirror function only on zSeries hosts. The z/OS Global Mirror function mirrors data on the storage unit to a remote location for disaster recovery. It protects data consistency across all volumes that you have defined for mirroring. The volumes can reside on several different storage units. The z/OS Global Mirror function can mirror the volumes over several thousand kilometers from the source site to the target recovery site.

With z/OS Global Mirror, you can suspend or resume service during an outage. You do not have to end your current data-copy session. You can suspend the session, and then restart it. Only data that changed during the outage must be resynchronized between the copies.

If you must change the default setting for an outage, you can do so through the DS Storage Manager. See the DS Storage Manager online help for more information.

IBM service representatives must identify and set up the DS8000 Copy Services servers before you can use the DS8000 Copy Services Web interface.

z/OS Metro/Global Mirror

Metro/Global Mirror is an optional function that is available with the z/OS remote mirror and copy feature and is supported on the DS8000.

The mirroring capability of z/OS Metro/Global Mirror uses z/OS Global Mirror to mirror data at a production site to a location that is a long distance away. This function also uses Metro Mirror to mirror data from the same primary devices at the production site to a location within a metropolitan area. The z/OS Metro/Global Mirror function enables a 3-site high availability and disaster recovery z/OS solution for even greater protection from unplanned outages.

Chapter 9. Migrating data to a storage unit

The DS8000 supports over 90 operating systems. You can migrate data to a storage unit from these host and operating system environments. The planning and methods of data migration vary by environment.

When you replace existing storage, you can partition the storage so that its virtual disks are similar in configuration to the disk drives that they are replacing. New configurations must be large enough to accommodate the existing data. Also, when data is moved, the data must be in a certain state, typically requiring that updates or changes cease while the movement occurs. Depending on the amount of data that you are moving and your migrating method, data could be unavailable for an extended period of time, perhaps several hours.

Your system administrator must select the data migration method that is the best compromise between efficiency and impact on the users of your system.

The following list is representative of the host and operating system environments supported by the DS8000:

- AIX
- OS/400
- iSeries (IBM i5/OS)
- zSeries (z/OS, z/OS.e, z/VM, VSE/ESA, UNIX, and TPF1)
- TRU64 UNIX
- Windows 2000 Server and Advanced Server with SP4
- Windows 2003
- Windows XP with SP1
- Linux for Intel

Data migration is a service offered through IBM Global Services. Contact your IBM Representative for more information.

How to select a data migration method

Your system administrator selects the data migration method that is the best compromise between efficiency and impact on the users of the system.

Most methods of data migration affect the everyday operation of a computer system. When data is moved, the data must be in a certain state, typically requiring that updates or changes cease while the movement occurs. Depending on the amount of data that you are moving and your migrating method, data could be unavailable for an extended period of time, perhaps several hours. The following factors might contribute to the migration time:

- Creating new logical volumes or file systems
- Modifying configuration files
- Receiving integrity checks

The following items are more than likely among the topics considered by your system administrator to determine the best method to use to migrate your data:

- Management software provides simple robust methods that you can generally use during production without disturbing users.

- The AIX logical volume manager (LVM) provides methods that you can use at any time without disrupting user access to the data. You might notice a small performance degradation, but this is preferable to shutting down databases or requiring users to log off the system.

Notes:

- AIX and HP-UX 11.xx ship with logical volume management (LVM) software as part of the base operating system. LVM provides complete control over all disks and file systems that exist on an AIX system. HP-UX has similar volume management software.
- Sun Microsystems has a basic volume management product called Solstice, which is available for the Solaris systems.
- Linux systems also use the LVM.
- Methods that use backup and restore procedures generally have the most impact on the system usage. They require that databases and file systems be in quiescent states to ensure a valid snapshot of the data.

Considerations for replacing existing storage

There are various factors to consider (remapping, reformatting of disk drives, partitioning) when replacing existing storage.

When you are replacing existing storage, partition the storage so that its virtual disks are similar in configuration to the disk drives that they are replacing. New configurations must be large enough to accommodate the existing data.

Note: You might want to take advantage of this opportunity to do some remapping. The allocation and distribution of data does not have to be a straight one-to-one relationship, although that is possible. For instance, you can take advantage of using a maximum of 255 logical subsystems whereas the prior limitation was 32 logical subsystems.

To provide integrity checking of data and to allow different operating systems to use the storage unit, IBM formats the disks to a 524-byte sector. The 524-byte sector replaces the traditional 512-byte sector of fixed-block architecture DDMs. The 524-byte sector format requires that you reformat all disks that you are migrating to the storage unit. Because reformatting erases any data that is already contained on the drives, you must back up the data to some form of temporary storage.

You can use these drives for either direct migration of data or as temporary storage while the existing drives are moved. If this is not possible, use a removable media device, such as a tape drive, to temporarily store the data while you reformat the drives.

Reformatting open systems host DDMs

To provide integrity-checking of data and to allow different operating systems to use the storage unit, IBM formats storage unit disks to a 524-byte sector.

The 524-byte sector replaces the traditional 512-byte sector of fixed-block architecture disk drive modules (DDMs). The 524-byte sector format requires that you reformat all disks that you are migrating to the storage unit. Because reformatting will erase any data that is already contained on the drives, you must back up the data to some form of temporary storage.

You can use the DDM for either direct migration of data or as temporary storage while the existing drives are moved. If this is not possible, use a removable media device such as a tape drive to temporarily store the data while you reformat the drives.

The AIX LVM uses a fixed-byte sector of 512-bytes. When used in an storage unit, the format is a fixed-byte sector of 524-bytes. The data portion of the sector remains at 512-bytes. AS/400 and iSeries headers use eight additional bytes. A 2-byte sequence number and a 2-byte longitudinal redundancy check (LRC) complete the extra bytes. The storage unit LIC uses the sequence number and LRC. The storage unit does not transfer them to the host system.

Most fixed-block disk architectures use a fixed-byte sector of 512-bytes. This includes most UNIX systems, including AIX. When used in an storage unit, the format of a DDM has a fixed-byte sector of 524-bytes. The 524-byte sector format enables the storage unit to connect to a wide range of host systems and share data between them.

The 524-byte sector on the DDM is independent of the sector size exposed to the host for a given device type emulation. Most LUNs for open-systems hosts are exposed with a 512-byte sector.

The IBM AS/400 and iSeries host systems use eight bytes at the start of the sector. UNIX host systems do not use these eight bytes when you attach them to the storage unit. The data portion of the sector remains at 512-bytes for all host systems. A 2-byte sequence number and a 2-byte LRC increase the size of the sector to 524-bytes.

Overview of the logical volume management software

The logical volume manager (LVM) software migrates data at a single logical volume level, a physical volume (DDM) level, or entire contents of a volume group. The LVM software provides complete control over all disks and file systems.

The primary tasks involved in migrating data using the LVM rely upon the use of LVM commands that affect the following functions:

- Copying
- Mirroring
- Migrating

Direct copy is another method to migrate data. While this method uses LVM, its main focus is on the use of the UNIX **find** command and the **cpio** command.

Copying a complete logical volume

Use this information as a guide for using the logical volume manager (LVM) **cplv** command to copy a complete logical volume.

The AIX LVM provides the **cplv** command for copying logical volumes within volume groups or to different volume groups. You can use this command to create a new logical volume while running the command, or you can overwrite an existing logical volume. The following two examples shows how to use the **cplv** command.

```
# cplv -v datavg -y newlv oldlv
# cplv -e existinglv oldlv
```

In the first example, the **cplv** command copies data from the existing logical volume `oldlv`. It creates a new logical volume that it calls `newlv` (`-y`) in the volume group `datavg` (`-v`). If you omit the `-v` option, the volume group to which the existing logical volume belongs receives the new logical volume. When the **cplv** command creates a new logical volume, it creates the new volume with exactly the same characteristics as the existing logical volume.

In the second example, the **cplv** command copies data from existing logical volume `oldlv` to existing logical volume `existinglv` (`-e`). When you use the `-e` option, it overwrites the existing target-logical volume with the data from the source-logical volume. When you use the `-e` option, it maintains the characteristics of the existing target-logical volume. Take care when you use this option.

The **cplv** command is a good method for copying or migrating a single logical volume. Sometimes, however, you might need to migrate all the data from a physical volume.

Using mirroring to migrate data

Use the following information as guide for how to use the mirroring function of logical volume manager (LVM) to migrate data. Mirroring can only be used on logical volumes. It cannot be used for physical volumes.

Mirroring is an LVM task that you perform only on logical volumes to migrate data. The following example shows how to create a mirror copy of a logical volume using the **mkivcopy** command:

```
# mkivcopy -e m -s y -k data1v 2 hdisk3 hdisk7
.
.
.
# splitlvcopy -y splitlv data1v 1
```

The **mkivcopy** command options specify the following values:

- (-e m)** To use minimum inter-disk allocation policy
- (-s y)** To strictly allocate mirror copies on separate physical volumes
- (-k)** To synchronize new copies immediately
- data1v** The name of the logical volume where you want to start a mirroring operation
- 2** The number of copies that you want to make of the data (a maximum of 3)
- hdisk3 and hdisk7** The physical volumes where the logical volume resides
- hdisk3** The physical volume that already holds the data of `data1v`
- hdisk7** The physical volume that will hold the mirror copy and where you want to move the data

The second example is similar to example 1. The **mkxlvcopy** command starts mirroring the `infxlv` logical volume.

```
# mkxlvcopy -e m -s y -k infxlv 2 hdisk4 hdisk10
.
.
.
# rmlxlvcopy infxlv 1 hdisk4
```

In this example, `hdisk4` is where the data for `infxlv` already exists, and `hdisk10` is where the mirror copy resides. It is also where you want to ultimately move the data.

Using the **migratepv** command

Use the following information as a guide for how to use the logical volume manager (LVM) **migratepv** command to migrate data that is associated with physical volumes.

You can use the LVM **migratepv** command to migrate data that is associated with physical volumes.

The following examples show how to use the **migratepv** command.

```
# migratepv hdisk1 hdisk2
# migratepv -l data1v hdisk3 hdisk9
```

In the example, all data migrates from `hdisk1` to `hdisk2`. The **migratepv** command updates all LVM references. From the time that the command completes, the LVM no longer uses `hdisk1` to access data that was previously stored there. As the data is physically moved, the target physical volume must have enough spare physical volumes to accommodate data from the source physical volumes. After this command completes, you can remove the source-physical volume from the volume group.

To avoid disrupting users, use the **migratepv** command while the system is active.

The **migratepv** command migrates data by performing the following actions:

- Creating a mirror of the logical volumes that you are moving
- Synchronizing the logical volumes
- Removing the original logical volume

You can use the **migratepv** command to move data from one physical volume to another physical volume within the same volume group. The following usage parameters apply to this command:

```
migratepv [-i] [-l LVname] SourcePV DestinationPV...
```

Note: You can specify more than one destination physical volume.

First, identify the source disk from which you want to migrate the data. Then, identify the target disk to which you want to migrate the data. You can only migrate to disks that are already in the `rootvg` volume group. To get a list of disks that are

already in the rootvg volume group, run the **lsvg -p rootvg** command. The following example displays output when you issue the **lsvg -p rootvg** command:

```
# lsvg -p rootvg
rootvg:
PV_NAME  PV STATE  TOTAL PPs  FREE PPs  FREE DISTRIBUTION
hdisk0   active    515        116        57..00..00..00..59
hdisk1   active    515        515        00..00..00..00..00
```

Now, determine the space that is currently in use on the disk that you want to migrate. This is the total physical partitions (PPs) value minus the free PPs value for the desired disk. In the preceding example, refer to hdisk0, which is using (515 - 116) PPs or 399 physical partitions.

Next, find a disk or disks that have the available space. In this case, hdisk1 has 515 free physical partitions, which is more than the required space of 399 physical partitions.

The only situation that is specific to rootvg is if the desired disk contains the boot image. Generally, this is the logical volume called hd5. From the following partial list of **lsvg -l rootvg**, this logical volume has a type of boot:

LV NAME	TYPE	Ps	PPs	PVs	LV STATE	MOUNT POINT
hd5	boot	1	1	1	closed/syncd	N/A

To determine if the boot image is on the disk that you want to migrate, run the **lslv -l** command. The following example displays the output when you issue the **lslv -l** command:

```
# lslv -l hd5
hd5:N/A
PV          COPIES          IN BAND          DISTRIBUTION
hdisk0      001:000:000      100%            001:000:000:000:000
```

In this case, the boot image resides on the disk that you want to migrate. You must first move the boot image with the **migratepv -l** command.

Using the direct copy method

Use the following information as a guide for how to use the direct copy method to migrate data.

There are times when you must use the direct copy method to migrate data. While this method uses the logical volume manager, the primary focus is on the use of the UNIX **find** command and the **cpio** command. The **find** command generates the list of files to be migrated. The **cpio** command migrates the files on the list.

The easiest way to produce the list of files is with the UNIX **find** command. Pipe its standard output to the standard input of the **cpio** command. The following output shows a typical example of using the **cpio** command to move data.

```
# mount /dev/lv00 /mnt
# cd /data
# find . -print | cpio -pdmuv /mnt
.
.
.
# umount /mnt
# umount /data
# mount /dev/lv00 /data
```

For the example above, the following considerations can apply:

- Assume that you made a file system on the /dev/lv00 logical volume. AIX LVM uses this file system to view part or all the virtual disks that the storage unit has made available to the system.
- Mount the logical volume on a temporary mount point, in this case /mnt.
- Change directories to the directory at the top of the file system that you want to move (cd/data).
- Use the **find** command to produce a list of file names, which a pipe (|) passes to the **cpio** command.
- Unmount both file systems and mount the new file system over the original mount point directory when the migration is complete.

You will probably be unable to use the volume management methods if the database uses volume serial numbers in its licensing code or validity checking. If the database uses licensing methods or validity checking, you might be able to only export the database from its old locations or import the database to its new location.

The database software provides the mechanism to move the data. This can take the form of a standard database backup and restore if it does not have any specific tools for moving data.

Using backup and restore methods to migrate data

Use the following information as a guide for how to use the backup and restore methods to migrate data.

In some cases, the only method that is available to transfer data is to back it up to tape and then restore it to the new disk. This method is slower because tape devices require that you move the data twice. However, if you are removing disks before you install the DS8000, the only way to move the data is with a tape device.

You can find a number of different archive utilities on UNIX systems. The **cpio** command can also create and read archives on tape devices. Use the **-o** (output) option to create archives, and use the **-i** (input) option to read and extract data from archives.

AIX provides a pair of commands: **backup** and **restore**. The **backup** command has two different methods of creating archives, either by file name or by the i-node. Use the **restore** command to read any archive that is created by the **backup** command. Solaris and HP-UX provide the **dump** and **restore** commands, which backup and restore data by the i-node.

The **tar** command is available on all UNIX systems that support the DS8000 and is another way to create and extract tape archives. Use the **c** option to create archives. Use the **x** option to extract files from archives.

Migrating data from an eServer i5 400 or iSeries host system to the DS8000

Use the following information as a guide to determine how to migrate data from an eServer i5 400 or iSeries host system to the DS8000.

When you create the eSeries i5 400 or iSeries volume, you can create it as a protected or unprotected volume. This creation enables you to see the protected and unprotected models. Unprotected models can be software mirrored.

DS8000 supports 2107 device types with model A81-A87 (unprotected) and A01-A07 (protected) and you need a level of OS/400 that supports these. The model numbers correspond to 8.56 GB (A01/A81), 17.54 GB (A02/A82), 35.16 GB (A05/A85), 36 GB (A03/A83), 70.56 GB (A84/A04), 141.12GB (A86/A06), and 282.25 GB (A87/A07).

eSeries i5 400 and iSeries hosts require a separate device address for each drive in the subsystem. The DS8000 meets this requirement by reporting unique addresses for each virtual drive that is defined to the eSeries i5 400 or iSeries host. 32 LUNS per iSeries FC adapter are supported by the on the OS/400 as well as multipathing to a storage unit.

Because the eSeries i5 400 or iSeries host does not support 7133 drawers, you have no existing 7133 drawers to reformat. If you have an existing disk subsystem that you are replacing with a DS8000, you must migrate the data to the DS8000. Use your existing host utilities for the data migration.

You can select from several methods to migrate data to the DS8000:

- You can use the logical add and remove functions.
- You can use save methods and restore methods with tape devices.

You can also use these methods if you remove an existing disk subsystem before you install the DS8000.

Note: From a DS8000 logical configuration viewpoint, all eSeries i5 400 and iSeries volumes are RAID 5 and are protected. When you create the eSeries i5 400 or iSeries volume, you can create it as a protected or unprotected volume. This creation enables you to see the protected and unprotected models. Unprotected models can be software mirrored.

Migrating data from an S/390 or zSeries host system to the DS8000 storage unit

You can use various methods to migrate data from an S/390 or zSeries host system to a DS8000 storage unit. The correct method depends on your environment and system.

The following table highlights a few of the data migration methods available:

Environment or Operating System	Data Migration Method
S/390 environment	IBM TotalStorage z/OS Global Mirror, Remote Mirror and Copy (when available)
zSeries environment	IBM TotalStorage z/OS Global Mirror, Remote Mirror and Copy (when available)

Environment or Operating System	Data Migration Method
Linux environment	Remote Mirror and Copy (when available)
z/OS operating system	<ul style="list-style-type: none"> • DFDSS (simplest method) • DFSMSHsm • IDCAMS EXPORT/IMPORT (VSAM) • IDCAMS REPRO (VSAM, SAM, BDAM) • IEBCOPY (partitioned data set (PDS), including load module libraries, and PDSE) • ICEGENER (SAM), part of DFSORT • IEBGENER (SAM) • Specialized database utilities for CICS, DB2, or IMS • Amdahl Transparent Data Migration Facility (TDMF), available from IBM as a data migration service
z/VM operating system	DFSMS/VM with the following functions: <ul style="list-style-type: none"> • DASD Dump Restore—Used to dump data from disk to tape, restore data from tape to disk, and copy data between like disk drive volumes. You cannot use DDR to copy data between disk devices with different track formats. • CMDISK—A DIRMAINT command used to move minidisks from any device types that VM supports to any other type. • COPYFILE—A CMS command used to copy files or minidisks between devices with the same or different track modes. • PTAPE—A Control Program (CP) command used to dump spool files to tape and to load files from tape to disk.
VSE/ESA operating system	<ul style="list-style-type: none"> • Backup and Restore dialogs • VSE fastcopy command to move volumes and files between devices with identical track formats. • VSE ditto command to copy files. • VSE power command to transfer the spool queue from one device to another. • VSE REPRO or EXPORT/IMPORT function to move any VSAM data set.

Data migration from a UNIX-based host system to a DS8000 storage unit

For UNIX-based host systems, you can use a variety of methods for copying or moving data from one disk drive to another.

The UNIX-based host sees the storage unit as one or more generic DDMMs. You can use the following methods to migrate data to the storage unit from open-system hosts that are running UNIX:

- Volume management methods
- Direct copy method
- Backup and restore methods
- Dump and restore commands as well as other commands for example the **tar** command

Migrating data to a DS8000 storage unit using volume management methods

The UNIX volume manager method for migrating data is very similar to the method that is used by AIX systems. It involves the use of commands or a combination of commands that copy and move the data.

The following methods are available for the UNIX volume manager. The method you choose depends on your operating environment.

The **cplv** method

Use the **cplv** command to make entire copies of logical volumes. You can issue the **cplv** command without disrupting access to data that is contained within the logical volume that it is copying. However, there is one drawback to using this command. Active logical volumes might not be copied when you use this command. You must ensure that your logical volumes are closed before you run this command.

To ensure consistency and limit the problems that are associated with this command, proceed with the following process:

1. Close your logical volumes either by unmounting the file system where they are built on, or by shutting down any database that has an open raw logical volume.
2. Run the **cplv** command.
3. Update the system configuration file, */etc/filesystems*, to include the relevant configuration data on the new logical volume and file system.
4. Issue the **fsck** command to run the file system integrity checker to ensure data consistency within the new logical volumes.

The **migratepv** method

This is an ideal method for moving data without disrupting users. Use the **migratepv** command to create a mirror of each logical volume that is contained within the physical volume. Issuing this command also synchronizes both copies. The benefit of this command is that you can run it on an active system.

You might notice some performance degradation due to the automatic nature of the creation and the synchronization of the copy. This process writes each physical partition and locks it from access by any other process. This can slow down access to the data, but it ensures data integrity.

The **mkivcopy**, **splitivcopy** method

This method is ideal for creating a copy of logical volumes. Use the **mkivcopy** command to ensure data integrity by creating a mirror copy of the data and synchronizing it automatically. Do not run the **splitivcopy** command or the **cplv** command on an active logical volume. If processes are updating the data while the split is taking place, the consistency of the data on both copies cannot be guaranteed. After you run the **splitivcopy** command, update the system configuration files. Include the relevant configuration data regarding the new logical volume and file system. Then issue the **fsck** to run the file system integrity checker to ensure data consistency within the new logical volume.

The **mkivcopy**, **rmivcopy** method

This method is ideal for migrating active logical volumes when a slight performance loss is acceptable. This method creates and removes mirrors automatically, ensuring data integrity.

Migrating data using the direct copy method

Use the following information as a guide for determining when to use the direct copy method to migrate data for a UNIX operating system.

This method suffers from the same drawback as using the **cplv** and **splitivcopy** commands on logical volumes. You can use this method on active file systems, but data consistency between the original and new copies cannot be guaranteed. Use the **cpio -p** command only on file systems that are in a quiescent state to guarantee consistency between the copy and the original. This will cause some disruption of service to users to occur.

Migrating data using backup and restore methods

Use the following information as a guide for determining when to use backup and restore methods to migrate data.

All backup and restore methods require that the file systems or logical volumes that you are backing up be in a quiescent state. Typically, you run backups after hours when there is minimal use of the system and you can shut down databases.

Some databases require that you export the data before you move it between file systems or logical volumes. You must then import the data on to the new file system or logical volume. You might have to reconfigure the database so that it points to the new data locations.

Migrating data from a Windows 2000 or Windows Server 2003-based host system

Use the following information as a guide for how to migrate data from a Windows 2000 or Windows Server 2003-based host system.

You can migrate data from Windows 2000-based host systems into an storage unit by mirroring the data and splitting the mirrors (similar to the method described for UNIX systems). Create a full backup of the data before attempting any data migration. Migrating data maintains file ownership, permissions, and the date and time while the copy is made. The drag-and-drop function and the **xcopy** command might cause ownership problems to date-related information in the copy of the new data.

Mirroring the data and splitting the mirrors requires some outage to rename drive letters at the end of the migration. If the mirroring tool that the operating system supplies cannot be used (for example, if drive sizes in the old and new systems are not the same), you might be able to use a third-party mirroring tool to accomplish the same goal. If an outage occurs, you can copy Windows 2000 data to the new disk using the **scopy** command.

Migrating data between storage units

Migrating data between storage units requires the use of Global Copy or IBM TotalStorage Metro Mirror. Perform this process to migrate data between storage units.

1. Establish paths between all Site A LSSs and Site B LSSs.
2. Establish all necessary volume pairs between the Site A storage units and the Site B storage units. You can use either the Global Copy or the IBM TotalStorage Metro Mirror feature.
3. Convert the pairs to IBM TotalStorage Metro Mirror after the bulk copy is complete if you use Global Copy.
4. Ensure that all remote mirror and copy volumes are in duplex state.
5. Quiesce the applications before you switch to the new volumes.
6. End all remote mirror and copy pairs across all storage units.
7. End all the remote mirror and copy paths from the Site A storage units to the Site B storage units.
8. Resume all applications that point to devices at the Site B storage units.

Moving workloads when migrating data between storage units

Moving workloads when migrating data between storage units requires the use of Global Copy or IBM TotalStorage Metro Mirror. Perform this process to move workloads between storage units.

1. Establish paths between all Site A LSSs and the Site B LSSs.
2. Establish all necessary volume pairs between the Site A storage units and the Site B storage units. You can use either Global Copy or IBM TotalStorage Metro Mirror.

If you use Global Copy, convert the pairs to IBM TotalStorage Metro Mirror after the bulk copy is complete.
3. Ensure that all remote mirror and copy volumes are in duplex state.
4. Quiesce all of the applications.
5. End all remote mirror and copy pairs across all ESSs.
6. End all remote mirror and copy paths.
7. Establish remote mirror and copy paths from the Site B storage units to the Site A storage units.
8. Establish all IBM TotalStorage Metro Mirror pairs from the Site B storage units to the Site A storage units.
9. Suspend all remote mirror and copy pairs. This causes the storage units to maintain a changed-track bitmap for all suspended remote mirror and copy primary volumes.
10. Resume all the Site A applications. The Site A applications use the Site B storage units disks.
11. Resume all the Site A applications that use the Site A storage units disks.

Offloading workloads when migrating data between storage units

Offloading workloads when migrating data between storage units requires the use of Global Copy or IBM TotalStorage Metro Mirror.

1. Establish paths between all Site A LSSs and Site B LSSs.
2. Establish all critical volume pairs between Site A storage units and Site B storage units, using either Global Copy or IBM TotalStorage Metro Mirror.

3. If you use Global Copy, convert the pairs to IBM TotalStorage Metro Mirror after the bulk copy is complete.
4. Ensure that all remote mirror and copy volumes are in duplex state at a point in time.
5. Issue database commands or file system commands to flush host buffers to disk, and quiesce the database or the file system.
6. Invoke FlashCopy for all volumes that are associated with the database.
7. Resume operations of the database or the file system.

Chapter 10. Disaster recovery using Copy Services

One of the main reasons for using Copy Services functions is in preparation for a possible disaster by backing up, copying, and mirroring your data at the production and recovery sites.

You can create a disaster recovery solution using any of the following functions. These functions help improve a disaster recovery process by allowing a much shorter recovery time with little or no data loss.

Note: In a remote mirror and copy environment, when two storage units are set up in two geographically distinct locations, we refer to the sites as Site A (the production site) and Site B (the recovery site).

Failover and failback operations

A failover is the process of temporarily switching production to a backup facility (normally your recovery site) following a planned outage, such as a scheduled maintenance period or an unplanned outage, such as a disaster. A failback operation is the process of returning production to its original location. These operations use remote mirror and copy functions to help reduce the time that is required to synchronize volumes after sites are switched during planned or unplanned outages.

Global Mirror

The Global Mirror function provides a two-site extended distance remote copy option for disaster recovery and backup solution for the zSeries and open systems environments. This solution is based on existing Global Copy and FlashCopy functions. With Global Mirror, the data that the host writes to the storage unit at the production site is asynchronously shadowed to the storage unit at the recovery site. Global Mirror operates over high-speed, Fibre Channel communication links and is designed to maintain a consistent and restartable copy of data at a recovery site that can be located at virtually unlimited distance from the production site.

Performing failover and subsequent failback operations

In the event of a planned outage, such as a scheduled maintenance or an unplanned outage, such as a hardware failure that disables the operation of your production site, you can failover to your recovery site to continue operations. After your production site is operational, you can perform a subsequent failback operation to move production back to its original location.

Moving production to Site B after planned outages (failover)

When you schedule a planned outage at your production site (Site A), you can switch production to your recovery site (Site B), allowing the processing of data to resume at Site B. This process is known as a failover recovery.

The storage units at both Site A and Site B must be functional and reachable.

In a disaster recovery environment, when two storage units are set up in two geographically distinct locations, we refer to the storage unit at the production or local site as Site A and the storage unit at the remote or recovery site as Site B.

For this scenario, assume that all I/O to Site A has ceased because of a planned outage, such as a scheduled maintenance. The failover operation is issued to the

storage unit that will become the primary. That is, production will failover to Site B during this outage, which will make the target volumes at Site B convert to source volumes and cause them to suspend. Your original source volumes at Site A remain in the state they were in at the time of the site switch. See Table 3 for an example of the implementation of failover and fallback operations.

Note: The failover recovery operation does not reverse the direction of a remote mirror and copy pair. It will change a target volume into a suspended source volume, while leaving the source volume in its current state.

The following assumptions are made for this scenario:

- Applications are updating the source volumes located at Site A.
- Paths are established from Site A to Site B.
- Volume pairs are in duplex state.

Perform the following steps using the DS Storage Manager. You can also use the DS CLI to perform Copy Services functions.

1. When the planned outage window is reached, quiesce applications to cease all write I/O from updating the source volumes. Quiescing your applications should occur as part of a planned outage and the outage should be brief. Depending on the host operating system, it might be necessary to unmount the source volumes.
2. Perform a failover recovery operation to Site B. After the failover operation has processed successfully, the volumes at Site B transition from target to source volumes.
3. Depending on your path design and when the source storage unit becomes available, create paths in the opposite direction from Site B to Site A. You will need the paths in the opposite direction if you want to transfer the updates back to Site A.
4. Depending on your operating system, it might be necessary to rescan your fibre-channel connected devices (to remove device objects for the site A volumes and recognize the new source volumes).
5. Mount your target volumes (now the new source volumes) on the target storage unit at Site B.
6. Start all applications. After the applications start, all write I/O operations to the source volumes are tracked. Depending on your plans regarding Site A, the volume pairs can remain suspended (if you want to do offline maintenance).
7. When your scheduled maintenance is complete, schedule a time for a fallback recovery operation to initiate the transfer of data back to Site A. This process resynchronizes the volumes at Site A with the volumes at Site B.

Note: Fallback recovery operations are usually used after a failover recovery has been issued to restart mirroring either in the reverse direction (remote site to local site) or original direction (local site to remote site).

Table 3. Failover and fallback implementation

Step	Operation	Management required to	Format of source volume and target volume	Format of source and target volume pair	Result: Site A	Result: Site B
1						

Table 3. Failover and failback implementation (continued)

Step	Operation	Management required to	Format of source volume and target volume	Format of source and target volume pair	Result: Site A	Result: Site B
Disaster at Site A	Failover	Site B	Volume B, Volume A	Volume B1 : Volume A1	Volume A1 -> Volume B1 (Suspended) The volume pair might display as full or pending duplex state if host write operations have stopped.	Volume B1 -> Volume A1 (Suspended)
2 (Site A volumes must be in suspended state)						
Return production to Site A	Failback	Site A	Volume A, Volume B	Volume A1 : Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1
3a (Site B volumes should be in a suspended state)						
Return to production (Site B) Note: If Site A is still not operational; production can continue at Site B.	Failback	Site B	Volume B, Volume A	Volume B1 : Volume A1	Volume B1 -> Volume A1	Volume B1 -> Volume A1
3b (prepare to return to production (Site A) from production (Site B))	Failover	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume B1 -> Volume A1 (Suspended state; the volume pair might display full or pending state if host write operations have stopped.)
3c (Site A volumes should be in a suspended state)						
Return to production - Site A	Failback	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1

Returning production to local site after planned outages

After restoring operations at Site B, you can schedule a failback operation to synchronize data and to enable production to resume at your original site, Site A.

Before you run a failback operation, you must create paths from Site B to Site A between the specific LSSs.

For this scenario, assume that Site A is operational and that connectivity from Site B to Site A is available. Use this procedure to restart your production environment

using the volumes from Site B. After the failover operation processes, the volumes at Site B are available to be mounted and used by a host. You can use the volumes from Site B to restart your application.

The failback operation resynchronizes the volumes in the following manner depending on the volume state:

- If a volume at Site A is in simplex state, all of the data for that volume is sent from Site B to Site A.
- If a volume at Site A is in full-duplex or suspended state and without changed tracks, only the modified data on the volume at Site B is sent to the volume at Site A.
- If a volume at Site A is in a suspended state but has tracks that have been modified, the volume at Site B will discover which tracks were modified at any site and send both tracks that were changed on Site A and the tracks marked at Site B from Site A to Site B.

The following assumptions are made for this scenario:

- Paths from Site B to Site A are created.
- Remote mirror and copy volume pairs are created. Site B volume is the source volume of the failback operation. This volume was initially the target volume of the relationship.

Perform the following steps using the DS Storage Manager. You can also use the DS CLI to perform Copy Services functions.

Note: For details on individuals tasks, see the related topics section.

1. Perform a failback operation at Site B. As a result, data that has been modified at Site B for the specified volumes will be sent from Site B to Site A. Data that was modified at Site A after the switch to Site B will be overwritten by the data from Site B.
2. Before returning to normal operation, quiesce your application (still updating volumes at Site B) to cease all write I/O from updating the source volumes. Depending on the host operating system, it might be necessary to unmount the source volumes.
3. Perform another failover recovery operation at Site A. This process converts the full-duplex target volumes at the Site A to suspended source volumes. The volumes at Site A start the change recording process while in failover mode.
4. Depending on your operating system, it might be necessary to rescan Fibre Channel devices and mount the new source volumes at Site A.
5. After the failover operation completes, you can restart the applications in normal production at Site A and check for consistency. When your applications are started, all write I/O to the new source volumes will be tracked by the bitmap that is created.
6. From Site A, create paths in the opposite direction from Site A to Site B.
7. From Site A, perform another failover and failback operation using the source volumes. This process resynchronizes the volumes between your Sites A and B. Start your applications at Site A and check for consistency.
8. Mount your volumes at Site A and start your applications on your primary server.

Moving production to Site B after unplanned outages (failover)

In this situation, you have lost access to your production site, Site A. You can switch production to your recovery site (Site B) allowing the processing of data to resume at Site B. This process is known as a failover recovery.

In a disaster recovery environment, when two storage units are set up in two geographically distinct locations, we refer to the storage unit at the production or local site as Site A and the storage unit at the remote or recovery site as Site B.

For this scenario, assume that an unexpected failure occurs at Site A, making it unavailable. Production must be moved to Site B.

The failover operation is issued to the storage unit that will become the primary. That is, production will failover to Site B during this outage, which will make the target volumes at Site B convert to source volumes and cause them to suspend. Your original source volumes at Site A remain in the state they were in at the time of the site switch. When Site A is available again, application I/O is switched back from Site B to Site A. See Table 4 for an example of the implementation of failover and failback operations.

Perform the following steps using the DS Storage Manager. You can also use the DS CLI to perform Copy Services functions.

- 1. Perform a failover recovery operation to Site B. After the failover operation has processed successfully, the volumes at Site B transition from target to source volumes.
- 2. Mount your target volumes on your server at Site B.
- 3. Start your applications on your server at Site B.
- 4. After Site A recovers, proceed with the following steps, which are the first taken to recover the volumes at Site A.
 - a. Create paths between LSSs at Site B to Site A to allow the volumes at Site A to be synchronized with the Site B volumes.
 - b. Delete any remote mirror and copy volume relationships that still exist from the source volumes.
 - c. Wait until the volumes are in full duplex state, then schedule a time to perform a failback recovery operation using the volumes at Site A. This process resynchronizes the volumes at Site A with the volumes at Site B.

Note: Failback recovery operations are usually used after a failover recovery has been issued to restart mirroring either in the reverse direction (remote site to local site) or original direction (local site to remote site).

Table 4. Failover and failback implementation

Step	Operation	Manage-ment required to	Format of source volume and target volume	Format of source and target volume pair	Result: Site A	Result: Site B
1						

Table 4. Failover and failback implementation (continued)

Step	Operation	Management required to	Format of source volume and target volume	Format of source and target volume pair	Result: Site A	Result: Site B
Disaster at Site A	Failover	Site B	Volume B, Volume A	Volume B1 : Volume A1	Volume A1 -> Volume B1 (Suspended) The volume pair might display as full or pending duplex state if host write operations have stopped.	Volume B1 -> Volume A1 (Suspended)
2 (Site A volumes must be in suspended state)						
Return production to Site A	Failback	Site A	Volume A, Volume B	Volume A1 : Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1
3a (Site B volumes should be in a suspended state)						
Return to production (Site B) Note: If Site A is still not operational; production can continue at Site B.	Failback	Site B	Volume B, Volume A	Volume B1 : Volume A1	Volume B1 -> Volume A1	Volume B1 -> Volume A1
3b (prepare to return to production (Site A) from production (Site B))	Failover	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume B1 -> Volume A1 (Suspended state; the volume pair might display full or pending state if host write operations have stopped.)
3c (Site A volumes should be in a suspended state)						
Return to production - Site A	Failback	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1

Using Global Mirror for disaster recovery

Global Mirror is a long-distance, two-site, remote copy solution that uses asynchronous technology. Global Mirror is based on existing Global Copy and FlashCopy functions.

The following tasks are illustrated:

- Setting up your environment to use Global Mirror
- Using Global Mirror for a planned failover
- Using Global Mirror for an unplanned failover
- Using recovery procedures with and without the fast reverse feature

Setting up your environment to use Global Mirror

Use this process to set up your environment to use Global Mirror for disaster recovery. Global Mirror asynchronously copies data from a host to a recovery site and maintains data on a storage unit at the recovery site.

You can configure a Global Mirror environment using two sites and three volumes on two or more storage units.

Note: The storage unit at the local site is designated as the Global Mirror master and all other local (or production) storage units are designated as subordinate storage units. The master sends commands to its subordinate storage units. These subordinates work together to create a consistency group and to communicate the FlashCopy commands to the recovery (or remote) site. All status is relayed back to the Global Mirror master.

When you create volume pairs using Global Copy, data on the volumes is not time consistent and generally not useful for the recovery site unless all of the data from the local site has been copied up to a specific point in time. To be useful for disaster recovery, the target volumes must periodically be made consistent and a copy of the data made. The following procedures describe how to set up your environment to use Global Mirror and create a consistent backup copy of your data.

1. Perform the initial planning for the use of Global Mirror. One of the first things the administrator of the site must do to use Global Mirror is determine the elements that are involved in the Global Mirror process. This includes the following:
 - The storage units that will participate in the session:
 - The storage unit that will be the Global Mirror master storage unit

Note: If you are using only one storage unit in your Global Mirror session, then that storage unit will be your master storage unit.

 - Other participating storage units that will be Global Mirror Subordinate storage units
 - The LSSs and volumes within those LSSs that are going to participate in the session:
 - Global Copy source volumes at the local sites
 - Global Copy target volumes (FlashCopy source volumes) at the recovery sites
 - FlashCopy target volumes at the recovery sites
 - LSSs to be used to connect the master and subordinate storage units
 - The following considerations are part of your management planning for this session, remembering that a session is a collection of volumes that are managed together to create consistent copies of data:
 - An LSS can only be assigned to one session.
 - A Global Mirror master storage unit coordinates the consistency group that is identified by a single session ID.
 - You are responsible for correctly defining sessions and managing the volumes within that session.
 - A maximum of eight storage units are allowed per session.
 - Data paths between Global Copy pairs
 - Global Mirror control paths between the master and subordinate storage units

Note: This is only necessary if you have subordinate storage units in your Global Mirror session.

2. Create fibre-channel paths between:
 - All Global Copy source and target pairs
 - Master and subordinate storage units
3. Create Global Copy pairs from the local storage units to the recovery storage units.

Note: The Global Copy source volumes are not active in the Global Mirror session until they have completed the first pass of their initial copy.

4. Create FlashCopy relationships at the recovery site between the Global Copy secondary volumes and the FlashCopy target volumes. Select all of the following options:

Enable Change Recording

Select this option to activate change recording on the volume pair participating in a FlashCopy relationship.

Note: The Persistent FlashCopy option is automatically selected with this option.

Inhibit writes to target volume

Select this option to ensure that writes are inhibited on the target volume. This ensures data consistency on the target volume. If you select the Inhibit writes to target option, the change recording feature is not active on the target volume, only on the source volume.

Note: Do not select the **Initiate background copy** option so that data will only be copied from the source volume to the target volume if a track on the source volume is modified.

5. Create your Global Mirror session.
6. Add volumes to your Global Mirror session.

Using Global Mirror for a planned failover and failback

Use this process to manage a planned failover involving two sites, which are referred to as local (or production) and remote (or recovery) sites, using Global Mirror.

During a planned outage, all storage units, servers, and networks are functional. Last and current states of all components are well-defined.

Examples for planned outages are system maintenance, disaster recovery tests, and training. The objective of the failover and failback procedures is to continue with consistent and current data (without any loss of transactions).

The procedure that is presented in this section makes the following assumptions:

- All storage units (local and remote) are functional and accessible.
- Applications are updating the primary volumes that are located at the local site.
- Paths are created from the local to the remote site.
- You have already created an environment to manage a planned outage.

Perform the following steps using the DS Storage Manager to begin the planned failover. You can also use the DS CLI.

Note: For details on individual tasks, see the related topics section.

1. Quiesce host I/O that is updating the primary volumes that is located in the local site.
2. Check the status of Global Mirror to ensure that Global Mirror is in a good state and that a consistency group has formed successfully at the desired interval.
3. Wait until two consistency groups form successfully.
4. Pause Global Mirror and suspend the Global Copy pairs (A->B).
5. Create paths from the B volumes to the A volumes.
6. Perform recovery failover on the B volumes (B->A).
7. Perform the fast reverse restore process from the C volumes to the B volumes, selecting the **Initiate background copy** option.

Note:

- a. When you initiate the fast reverse restore process, Volume C becomes unusable.
 - b. There must be *no* I/O allowed to the B or C volumes during the fast reverse restore process.
8. Wait for the background copy to complete before continuing to the next step. The C to B FlashCopy relationship ends when the background copy completes.
 9. Initiate the FlashCopy from Volume B back to Volume C. Ensure that you also select the **Enable Change Recording** and **Inhibit writes to target volume** options. This creates a backup copy of the consistency group before applications begin to update the B volumes.
 10. Start the host I/O at the remote site on the B volumes. You will remain in production on the remote site in this configuration until you are ready to return production to the local site.
 11. When you are ready to return production to the local site, perform recovery failback (B->A) to resynchronize the A volumes. The application at the remote site remains active.
 12. After the resynchronize process has completed its first pass, quiesce the applications at your remote site so that the resynchronization can complete.
 13. When the resynchronization has completed (no out-of-sync tracks), perform recovery failover and failback with Global Copy on Volume A to re-create the Volume A to Volume B Metro Mirror relationship.
 14. Start the host I/O at the local site on the A volumes.
 15. Resume the Global Mirror process.

Using Global Mirror for an unplanned failover and failback

Use this process to manage an unplanned failover and failback that involves two sites which are referred to as the local (or production) and remote (or recovery) sites.

Global Mirror provides two-site extended distance remote copy disaster recovery. When a disaster occurs at the local site, you must initiate the failover and failback recovery of consistent data on the remote site. Host activity can resume on the local site when the host recovers but not before a consistent set of data is copied to all primary volumes on the local site.

With Global Mirror, the data that the host writes to the storage unit at the local site is asynchronously shadowed to the storage unit at the remote site. A consistent copy of the data is then automatically maintained on the storage unit at the remote site.

Use of Global Mirror does not guarantee against data loss. During a disaster, data can only be restored to the last known consistent increment that was created. This means that data that is written to the primary site and is waiting to be transferred to the secondary site is lost whenever the two storage units can no longer communicate.

The following considerations apply when you use the Global Mirror recovery process:

- The Global Mirror master might still be running at the local site, especially if the disaster at the local site is a rolling disaster in which not all components fail simultaneously.
- The consistent copy at the remote site is *not* the secondary volume, but it is the FlashCopy target whose source is the secondary volume.
- Formation of a consistency group might have been in progress at the time of the failure.
- You can speed up recovery processing by choosing the **Fast Reverse** restore process that is explained later in this section.

Perform the following steps to use Global Mirror for an unplanned failover using the DS Storage Manager. You can also use the DS CLI for these steps.

Note: For details on individual tasks, see the related topics section.

1. Enter the Global Mirror session at the local site.

Note: Wait until the master storage unit completes the termination processing or enters the fatal state before continuing with the next step. Of course, this might not be possible if the local site has completely failed. If that is the case, proceed to the next step without waiting.

2. Issue a recovery failover request on the Global Copy volumes pair to force a stop of the Volume A to Volume B extended distance relationship and create a Volume B to Volume A Global Copy relationship.

Note: All B volumes must successfully process the recovery failover request before you can move to the next step.

3. Look at the session properties for Volume B and Volume C to ascertain the state of the consistency group between the B volume and C volume. You are looking primarily at the FlashCopy relationships and your analysis determines your next step in the recovery process. Act on your analysis as follows:
 - a. FlashCopy relationships are nonrestorable and all the sequence numbers are equal. No action to the consistency group is necessary.
 - b. FlashCopy relationships are restorable and all the sequence numbers are equal. Issue the FlashCopy **Discard changes** command to all the FlashCopy relationships in the consistency group.
 - c. All the FlashCopy sequence numbers are equal and at least one of the FlashCopy relationships is nonrestorable. Issue a FlashCopy **Commit changes** command to all the FlashCopy relationships in the consistency group that are restorable.

- d. You have a mixed list of FlashCopy relationships, some are restorable and some are nonrestorable. The sequence numbers of the relationships that are restorable all have the same sequence number. The sequence numbers of the relationships that are nonrestorable are all equal, but they have a different number from the sequence number of those that are restorable. Issue a FlashCopy **Commit changes** command to all the FlashCopy relationships in the consistency group that are restorable.
- e. You have a mixed list of FlashCopy relationships, some are restorable and some are nonrestorable. The sequence numbers are not the same within each type of relationship. The recovery plan cannot continue. The Global Mirror process has been corrupted. If the Global Mirror process has been corrupted you must recover your data using your last good backup.

Note: When the state of all the FlashCopy relationships are known, you might want to initiate a tape backup of Volume C.

4. Issue the fast reverse restore process from the C volumes to the B volumes, selecting the **Initiate background copy** option.

Note:

- a. When you initiate the fast reverse restore process, Volume C becomes unusable.
 - b. There must be *no* I/O allowed to the B or C volumes during the fast reverse restore process.
 - c. If you do not want to use the fast reverse restore process, use the Recovering from a disaster without using fast reverse restore procedure instead of this step.
5. Wait for the background copy to complete before continuing to the next step. The C to B FlashCopy relationship ends when the background copy completes.
 6. Initiate the FlashCopy from Volume B back to Volume C. Ensure that you also select the **Enable Change Recording** and **Inhibit writes to target volume** options. This creates a backup copy of the consistency group before applications begin to update the B volumes.
 7. Start the host I/O at the remote site on the B volumes. You will remain in production on the remote site in this configuration until you are ready to return production to the local site.
 8. When you are ready to return production to the local site, perform the recovery failback (B->A) to resynchronize the A volumes. The application at the remote site remains active.
 9. After the resynchronize process has completed its first pass, quiesce the applications at your remote site so that the resynchronization can complete.
 10. When the resynchronization has completed (no out-of-sync tracks), perform recovery failover and failback with Global Copy on Volume A to re-create the Volume A to Volume B Metro Mirror relationship.
 11. Start the host I/O at the local site on the A volumes.
 12. Resume the Global Mirror process.

Recovering with and without the fast reverse restore feature

The fast reverse restore feature, when used with Global Mirror functions, allows a FlashCopy relationship to be reversed without waiting for the background copy of the original FlashCopy to complete.

You can reverse a FlashCopy relationship that has change recording enabled and restore the tracks from the target volume to the source volume. However, the background copy process of the original FlashCopy must complete before you can reverse the order of the FlashCopy relationship to its original source and target relationship.

With the fast reverse restore feature, you can reverse a FlashCopy relationship that has change recording enabled without waiting for the background copy of the original FlashCopy to complete. It is not required that you use the fast reverse restore feature during disaster recovery operations; however, the time to recover could be significantly longer if you do not.

Recovering from a disaster using the fast reverse process

Use this process for the Global Mirror fast reverse restore feature, which allows a FlashCopy relationship to be reversed without waiting for the background copy of the original FlashCopy to complete.

You can use the fast reverse restore feature of FlashCopy to speed up the recovery process. Using the DS Storage Manager, you make the following selections:

1. Create a FlashCopy target on the existing Metro Mirror source by choosing the FlashCopy target and a source Metro Mirror .
2. Select the **Fast Reverse** restore option.

FlashCopy copies the partial data that had been copied to Volume C before the disaster occurred to Volume B. This creates a consistent copy on Volume B. Processing also provides a background copy of the consistent group for Volume B.

Note: When you initiate the fast reverse restore process, Volume C becomes unusable.

Recovering from a disaster without the fast reverse process

Use this procedure in place of the fast reverse process. The fast reverse process allows a FlashCopy relationship to be reversed without waiting for the background copy of the original FlashCopy to complete.

It is not necessary to use the fast reverse feature of FlashCopy for your recovery. However, the wait for recovery is significantly longer if you do not.

1. Initiate a background copy on the Volume B to Volume C FlashCopy relationship.
2. Wait for the background copy to complete. Observe the out-of-sync tracks.
3. Select the **Reverse relationship** and **Initiate background copy** options on the Volume B to Volume C FlashCopy relationship. This reverses the FlashCopy relationship such that Volume C is the source and Volume B is the target.

Appendix. Troubleshooting

The topics in this section provide troubleshooting information related to your DS8000. Topics covered include analyzing, verifying, understanding, managing, and handling various problems.

Analyzing generic alert traps

The storage unit generates generic alert traps. The traps are sent by an SNMP agent to report asynchronous events.

Generic trap alert identifiers

coldStart (Generic 0)

Issued whenever the SNMP agent is reinitializing. Configuration data has changed.

warmStart (Generic 1)

Issued whenever the SNMP agent is reinitializing. Configuration data might change.

authenticationFailure (Generic 4)

Issued whenever an SNMP message was received but could not be authenticated.

Analyzing service information messages (SIMs) for S/390 and zSeries systems

Service information messages (SIMs) are generated by a storage unit for S/390 and zSeries hosts. Before installation, you can use the customization work sheets to record the following service information: the SIM severity level and how often the storage unit sends the SIMs (0-9 times) to the console. During installation, either you or the IBM service sales representative must enter this information into the system.

SIM message types

The following SIM message types are generated by the storage unit.

Direct access storage device (DASD) SIM

Tracks DDM failures and problems.

Media SIM

Tracks data check problems on the media.

Storage unit SIM

Tracks storage unit, power control, and other hardware problems.

SIM severity levels

1 acute

An irrecoverable error with possible loss of data. This severity level only applies to DASD SIMs.

2 serious

An irrecoverable error or a data check with loss of access to data.

3 moderate

A system path is not operational and performance might be degraded. This severity level does not apply to media SIMs.

4 service

A recoverable error, equipment checks, or data checks. You can defer repair.

Understanding problem status designations

Each storage unit problem creates a serviceable event. The serviceable event status is either Open or Closed.

Problem status designations

The designations explain where the problem resides in the resolution process.

Open A problem has occurred that requires service. The status will remain in open prior to and during the repair.

Closed

The repair was completed, and that changed the status to closed.

Managing product-specific alert traps

Product-specific alert traps provide information about problems that the storage unit detects and that require corrective action. Either you or the service provider must perform some action for each of these problems. The storage unit generates the following generic, trap-6, product-specific alert traps: 1, 2, 3, 100, 101, 102, 200, 202, 210, 211, 213, 214, 215, 216, 217.

Specific information provided for trap-6, alert trap 1

Each generic, trap-6, product specific alert trap 1 provides the following information.

Problem ID

The identifier assigned by the storage unit for the particular problem that caused the trap.

Presenting time

The time and date of the most recent problem report.

Description

The description of the error condition for the problem ID.

User action

The recommended actions that you need to take to resolve this error condition.

Product-specific alert traps 1 and 2

Product-specific alert traps contain the identifier that the storage unit assigns for the particular problem that caused the trap. The character string *Problemid=n* (where *n* is the problem identifier) is in the description field of the trap. The function for generic, trap-6, product-specific trap 1, and trap 2 on the storage unit is the same. Product-specific alert trap 1 is set as the default.

Product-specific alert traps 1 and 2 generate the following information:

yyyy/mm/dd hh:mm:ss zzz

Provides the time and date of the most recent problem occurrence. (zzz represents the time zone.)

data Provides the following detailed information about the trap:

- Storage unit machine type
- Model
- Serial number
- Problem ID
- Severity

partnumber

Provides the failure resource unit (FRU) and FRU part number.

Product-specific alert traps 3, 100, 101, 102, 200, 202, 210, 211, 212, 213, 214, 215, 216, and 217

These alert traps generate the following information:

yyyy/mm/dd hh:mm:ss zzz

Provides the time and date of the most recent problem occurrence. (zzz represents the time zone.)

description

Provides a short textual description of the trap.

data Provides detailed information about the trap.

Obtaining services for a storage unit

The following are IBM services that you can obtain to benefit the processing associated with your storage unit.

Hardware problems

The storage unit is capable of remote error notification and remote support for those machines that are under warranty or a maintenance agreement. An IBM service support representative (SSR) configures your storage unit for remote service during installation.

Data migration

IBM provides a service through Global Services to help you with your data migration needs. Contact your IBM representative for more details.

Command Line Interface (CLI)

IBM provides a service through Global Services to help you with using the DS8000 CLI in your system environment. Contact your IBM representative for more details.

Analyzing normal operation problems

Try the following actions to resolve a problem.

When the storage unit encounters a error that requires action, it illuminates one of the message lights on the expansion enclosure. Also informational messages are issued through the storage unit when special events occur.

If your warranty covers the storage unit or you have a service maintenance agreement, you can respond in the following way:

- Contact your authorized service representative.
- Describe the error message, error code, or problem that you have observed.

Note: Your warranty agreement or service maintenance agreement might be affected if you perform your own maintenance.

Use the following topical descriptions of the messages generated by or through the storage unit to help you describe the problem to your service representative.

Managing informational messages from the storage unit

Informational messages are issued through the storage unit as special events occur. Your system administrator determines how these messages should be handled.

Type of informational message

An informational message is issued when your service provider runs the customer-notification diagnostic test. This test verifies that e-mail messages are being received by those who should receive them. You need to attach your LAN to the DS Network to receive e-mail messages from the storage unit.

Managing the storage unit error messages

The storage unit generates error messages when it detects a situation that requires customer action. This section describes the type of information provided so that you can take the necessary steps to resolve the error condition. In most cases you will need to call your IBM service representative. By providing the listed information your IBM service representative will have an idea where to begin to resolve the problem.

Purpose

The error messages from the storage unit typically contain the following fields.

Product manufacturer ID and date

The ID of the storage unit and the date that it was manufactured.

Storage unit location

The installer enters the storage unit location during the initial installation of the product.

Product machine type and model number

Assigned by IBM at time of manufacturing.

Product serial number

Assigned by IBM at time of manufacturing.

Customer voice phone number

The phone number for customer voice contact.

LMC level of local storage server

The level of the licensed machine code (LMC) of your primary storage unit.

LMC level of remote storage server

The level of the LMC of your secondary or backup storage unit.

Report time and date stamp

The time that this report was generated.

Problem ID

The problem ID that is assigned to this problem by the storage unit. The service provider uses this problem ID to access detailed problem information.

SRN/SRC

A detailed error code that the service provider uses.

Problem status

The problem status state.

Description

A description of the problem.

Additional message

Any additional information that is available.

Reporting resource

The coded resource name that the service provider uses during the repair process.

Failure occurred

The date and time when the failure first occurred.

Last occurrence

The date and time the last occurrence was noted.

Failure count

The number of times that this failure occurred.

Presentation interval

The time between successive e-mail copies of this problem.

Remaining presentations

The number of additional times this e-mail notification will be sent.

Isolation procedure

A pointer to a special procedure in the online service information center.

Failure actions

Actions that the service provider can take.

Probable cause

Information for the service provider.

Failure cause

Information for the service provider.

The following fields are the most useful to you in identifying DDM failures:

- Description
- Reporting resource
- Last occurrence

Accessibility

Accessibility features provide users who have disabilities with the ability to successfully access information and use technology.

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features

These are the major accessibility features in the IBM TotalStorage DS8000 information:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. IBM Home Page Reader version 3.0 has been tested.
- You can operate features using the keyboard instead of the mouse.

Navigating by keyboard

You can use keys or key combinations to perform operations and initiate menu actions that can also be done through mouse actions. You can navigate the IBM TotalStorage DS8000 information from the keyboard by using the shortcut keys for your browser or Home Page Reader. See your browser Help for a list of shortcut keys that it supports. See the following Web site for a list of shortcut keys supported by Home Page Reader: http://www-306.ibm.com/able/solution_offerings/keyshort.html

Accessing the publications

You can find HTML versions of the IBM TotalStorage DS8000 information at the following Web site:
<http://www.ehone.ibm.com/public/applications/publications/cgibin/pbi.cgi>

You can access the information using IBM Home Page Reader 3.0.

Notices

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "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.

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 measurement 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.

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.

Terms and conditions for downloading and printing publications

Permissions for the use of the information you have selected for download are granted subject to the following terms and conditions and your indication of acceptance thereof.

Personal Use: You may reproduce this information for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of this information, or any portion thereof, without the express consent of the manufacturer.

Commercial Use: You may reproduce, distribute and display this information solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of this information, or reproduce, distribute or display this information or any portion thereof outside your enterprise, without the express consent of the manufacturer.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the information or any data, software or other intellectual property contained therein.

The manufacturer reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the information is detrimental to its interest or, as determined by the manufacturer, the above instructions are not properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations. THE MANUFACTURER MAKES NO GUARANTEE ABOUT THE CONTENT OF THIS INFORMATION. THE INFORMATION IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

All material copyrighted by IBM Corporation.

By downloading or printing information from this site, you have indicated your agreement with these terms and conditions.

Trademarks

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

- AIX
- DB2
- DFSMS/MVS
- DFSMS/VM
- e (logo)
- Enterprise Storage Server
- ES/9000
- ESCON
- FICON
- FlashCopy
- HACMP
- IBM
- Intellistation
- MVS/ESA
- Netfinity
- NetVista
- Operating System/400
- OS/400
- RS/6000
- S/390
- Seascape
- SNAP/SHOT
- SP
- System/390
- TotalStorage
- Versatile Storage Server
- Virtualization Engine
- VSE/ESA
- z/Architecture
- z/OS
- z/VM
- zSeries

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

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

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

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

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

Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.

Federal Communications Commission (FCC) statement

This equipment has been tested and complies with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the users authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Industry Canada compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conform à la norme NMB-003 du Canada.

European community compliance statement

This product is in conformity with the protection requirements of EC Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

Germany only

Zulassungsbescheinigung laut Gesetz ueber die elektromagnetische Verträglichkeit von Geraeten (EMVG) vom 30. August 1995.

Dieses Geraet ist berechtigt, in Uebereinstimmung mit dem deutschen EMVG das EG-Konformitaetszeichen - CE - zu fuehren.

Der Aussteller der Konformitaetserklaeung ist die IBM Deutschland.

Informationen in Hinsicht EMVG Paragraph 3 Abs. (2) 2:

Das Geraet erfuehlt die Schutzanforderungen nach EN 50082-1 und EN 55022 Klasse A.

EN 55022 Klasse A Geraete beduerfen folgender Hinweise:

Nach dem EMVG:

"Geraete duerfen an Orten, fuer die sie nicht ausreichend entstoert sind, nur mit besonderer Genehmigung des Bundesministeriums fuer Post und Telekommunikation oder des Bundesamtes fuer Post und Telekommunikation betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Stoerungen zu erwarten sind." (Auszug aus dem EMVG, Paragraph 3, Abs.4)

Dieses Genehmigungsverfahren ist nach Paragraph 9 EMVG in Verbindung mit der entsprechenden Kostenverordnung (Amtsblatt 14/93) kostenpflichtig.

Nach der EN 55022:

"Dies ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funkstoerungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Massnahmen durchzufuehren und dafuer aufzukommen."

Anmerkung:

Um die Einhaltung des EMVG sicherzustellen, sind die Geraete wie in den Handbuechern angegeben zu installieren und zu betreiben.

Japanese Voluntary Control Council for Interference (VCCI) class A statement

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Korean Ministry of Information and Communication (MIC) statement

Please note that this device has been certified for business use with regard to electromagnetic interference. If you find this is not suitable for your use, you may exchange it for one of residential use.

Taiwan class A compliance statement

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

VS07171L

Glossary

This glossary includes terms for the IBM TotalStorage and other Resiliency Family products.

This glossary includes selected terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172–1990, copyright 1990 by the American National Standards Institute (ANSI), 11 West 42nd Street, New York, New York 10036. Definitions derived from this book have the symbol (A) after the definition.
- *IBM Terminology*, which is available online at the following Web site: <http://www-306.ibm.com/ibm/terminology/index.html>. Definitions derived from this source have the symbol (GC) after the definition.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions derived from this book have the symbol (I) after the definition. Definitions taken from draft international standards, committee drafts, and working papers that the ISO/IEC JTC1/SC1 is developing have the symbol (T) after the definition, indicating that final agreement has not been reached among the participating National Bodies of SC1.

This glossary uses the following cross-reference forms:

- See** Refers the reader to one of two kinds of related information:
- A term that is the expanded form of an abbreviation or acronym. This expanded form of the term contains the full definition.
 - A synonym or more preferred term

See also

Refers the reader to one or more related terms.

Contrast with

Refers the reader to a term that has an opposite or substantively different meaning.

Numerics

- 750** A model of the Enterprise Storage Server featuring a 2-way processor with limited physical storage capacity. This model can be updated to the model 800.
- 800** A model of the Enterprise Storage Server featuring a standard processor or an optional Turbo processor. The Model 800 supports RAID 5, RAID 10, and 15000 rpm drives. Model 800 supersedes Model F20.
- 2105** The machine number for the IBM TotalStorage Enterprise Storage Server. Models of the Enterprise Storage Server are expressed as the number 2105 followed by "Model <xxx>", such as 2105 Model 800. The 2105 Model 100 is an Enterprise Storage Server expansion enclosure that is typically referred to simply as the Model 100.
- 2107** The machine type for the IBM TotalStorage DS8000 series. Models for the DS8000 series include the base units 921, 922 and 9A2 and expansion units 92E and 9AE.
- 3390** The machine number of an IBM disk storage system. The Enterprise Storage Server, when interfaced to IBM zSeries hosts, is set up to appear as one or more 3390 devices, with a choice of 3390-2, 3390-3, or 3390-9 track formats.
- 3990** The machine number of an IBM control unit.
- 7133** The machine number of an IBM disk storage system. The Model D40 and 020 drawers of the 7133 can be installed in the 2105-100 expansion enclosure of the ESS.

A

access

- 1) To obtain the use of a computer resource.

2) In computer security, a specific type of interaction between a subject and an object that results in flow of information from one to the other.

access-any mode

One of the two access modes that can be set for the storage unit during initial configuration. It enables all fibre-channel-attached host systems with no defined access profile to access all logical volumes on the storage unit. With a profile defined in DS Storage Manager for a particular host, that host has access only to volumes that are assigned to the WWPN for that host. See also *pseudo host* and *worldwide port name*.

ACK See *request for acknowledgment and acknowledgment*.

agent A program that automatically performs some service without user intervention or on a regular schedule. See also *subagent*.

alert A message or log that a storage unit generates as the result of error event collection and analysis. An alert indicates that a service action is required.

allegiance

For zSeries, a relationship that is created between a device and one or more channel paths during the processing of certain conditions. See also *implicit allegiance*, *contingent allegiance*, and *reserved allegiance*.

allocated storage

The space that is allocated to volumes but not yet assigned. Contrast with *assigned storage*.

American National Standards Institute (ANSI)

An organization of producers, consumers, and general interest groups that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States. (A)

Anonymous

In the DS Storage Manager, the label on an icon that represents all connections that are using fibre-channel adapters between the storage unit and hosts but are not completely defined to the storage unit. See also *anonymous host*, *pseudo host*, and *access-any mode*.

anonymous host

Synonym for *pseudo host*. Contrast with *Anonymous* and *pseudo host*.

ANSI See *American National Standards Institute*.

APAR See *authorized program analysis report*. (GC)

API See *application programming interface*.

application programming interface

An interface that allows an application program that is written in a high-level language to use specific data or functions of the operating system or another program.

arbitrated loop

A fibre-channel topology that enables the interconnection of a set of nodes. See also *point-to-point connection* and *switched fabric*.

array An ordered collection, or group, of physical devices (disk drive modules) that is used to define logical volumes or devices. In the storage unit, an array is a group of disks that the user designates to be managed by the RAID technique. See also *redundant array of independent disks*.

ASCII (American National Standard Code for Information Interchange) The standard code, using a coded character set consisting of 7-bit coded characters (8 bits including parity check), that is used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. (A) Some organizations, including IBM, have used the parity bit to expand the basic code set.

assigned storage

The space that is allocated to a volume and that is assigned to a port.

authorized program analysis report (APAR)

A request for correction of a defect in a current release of an IBM-supplied program. (GC)

availability

The degree to which a system or resource is capable of performing its normal function. See *data availability*.

B

bay The physical space that is used for installing SCSI, ESCON, and fibre-channel host adapter cards. The DS8000 has four bays, two in each cluster. See also *service boundary*.

bit 1) Either of the digits 0 or 1 when used in the binary numeration system. (T)
2) The storage medium required to store a single binary digit. See also *byte*.

block A string of data elements recorded or transmitted as a unit. The elements may be characters, words, or physical records. (T)

A group of consecutive bytes used as the basic storage unit in fixed-block architecture (FBA). All blocks on the storage device are the same size (fixed size). See also *fixed-block architecture* and *data record*.

byte 1) A group of eight adjacent binary digits that represent one EBCDIC character.
2) The storage medium required to store eight bits. See also *bit*.

C

cache A special-purpose buffer storage, smaller and faster than main storage, used to hold a copy of instructions and data obtained from main storage and likely to be needed next by the processor. (T)

cache fast write

A form of the fast-write operation in which the storage server writes the data directly to cache, where it is available for later destaging.

cache hit

An event that occurs when a read operation is sent to the cluster, and the requested data is found in cache. Contrast with *cache miss*.

cache memory

Memory, typically volatile memory, that a storage server uses to improve access times to instructions or data. The cache memory is typically smaller and faster than the primary memory or storage medium. In addition to residing in cache memory, the same data also resides on the storage devices in the storage unit.

cache miss

An event that occurs when a read operation is sent to the cluster, but the data is not found in cache. Contrast with *cache hit*.

call home

A communication link established between the DS8000 and a service provider. The DS8000 can use this link to place a call to IBM or to another service provider when it requires service. With access to the machine, service personnel can perform service tasks, such as viewing error logs and problem logs or initiating trace and dump retrievals. See also *heartbeat* and *remote technical assistance information network*.

cascading

1) Connecting network controllers to each other in a succession of levels to concentrate many more lines than a single level permits.

2) In high-availability cluster multiprocessing (HACMP), cascading pertains to a cluster configuration in which the cluster node with the highest priority for a particular resource acquires the resource if the primary node fails. The cluster node relinquishes the resource to the primary node upon reintegration of the primary node into the cluster.

catcher

A server that service personnel use to collect and retain status data that an DS8000 sends to it.

CCR See *channel command retry*.

CCW See *channel command word*.

CD See *compact disc*.

CEC See *computer-electronic complex*.

channel

For zSeries, the part of a channel subsystem that manages a single I/O interface between a channel subsystem and a set of control units.

channel command retry (CCR)

For zSeries, the protocol used between a channel and a control unit that enables the control unit to request that the channel reissue the current command.

channel command word (CCW)

For zSeries, a data structure that specifies an I/O operation to the channel subsystem.

channel path

For zSeries, the interconnection between a channel and its associated control units.

channel subsystem

For zSeries, the part of a host computer that manages I/O communication between the program and any attached control units.

channel-subsystem image

For zSeries, the logical functions that a system requires to perform the function of a channel subsystem. With ESCON multiple image facility (EMIF), one channel subsystem image exists in the channel subsystem for each logical partition (LPAR). Each image appears to be an independent channel subsystem program, but all images share a common set of hardware facilities.

CKD See *count key data*.

CLI See *command-line interface*. See also *Copy Services command-line interface*.

cluster

1) A partition capable of performing all DS8000 functions. With two clusters in the DS8000, any operational cluster can take over the processing of a failing cluster.

2) In the AIX operating system, a group of nodes within a complex.

cluster processor complex (CPC)

The unit within a cluster that provides the management function for the DS8000. It consists of cluster processors, cluster memory, and related logic.

command-line interface (CLI)

An interface provided by an operating system that defines a set of commands and enables a user (or a script-like language) to issue these commands by typing text in response to the command prompt (for example, DOS commands or UNIX shell commands). See also *Copy Services command-line interface*.

compact disc

An optically read disc, typically storing approximately 660 MB. CD-ROM (compact disc read-only memory) refers to

the read-only format used to distribute DS8000 code and documentation.

compression

1) The process of eliminating gaps, empty fields, redundancies, and unnecessary data to shorten the length of records or blocks.

2) Any encoding that reduces the number of bits used to represent a given message or record. (GC)

computer-electronic complex (CEC)

The set of hardware facilities associated with a host computer.

concurrent copy

A facility on a storage server that enables a program to make a backup of a data set while the logical volume remains available for subsequent processing. The data in the backup copy is frozen at the point in time that the server responds to the request.

concurrent installation of licensed internal code

Process of installing licensed internal code on a DS8000 while applications continue to run.

concurrent maintenance

Service that is performed on a unit while it is operational.

concurrent media maintenance

Service performed on a disk drive module (DDM) without losing access to the data.

configure

In storage, to define the logical and physical configuration of the input/output (I/O) subsystem through the user interface that the storage unit provides for this function.

consistency group

A group of volumes participating in FlashCopy relationships in a logical subsystem, across logical subsystems, or across multiple storage units that must be kept in a consistent state to ensure data integrity.

consistency group interval time

The value in seconds that indicates the length of time between the formation of consistency groups.

consistent copy

A copy of a data entity (a logical volume,

for example) that contains the contents of the entire data entity at a single instant in time.

console

A user interface to a server, for example, the interface provided on a personal computer. See also *IBM TotalStorage Management Console*.

contingent allegiance

For zSeries, a relationship that is created in a control unit between a device and a channel when the channel accepts unit-check status. The allegiance causes the control unit to guarantee access; the control unit does not present the busy status to the device. The allegiance enables the channel to retrieve sense data that is associated with the unit-check status on the channel path associated with the allegiance.

control path

The route that is established from the master storage unit to the subordinate storage unit when more than one storage unit participates in a Global Mirror session. If there is only one storage unit (the master) in the Global Mirror session, no control path is required.

control unit (CU)

- 1) A device that coordinates and controls the operation of one or more input/output devices, and synchronizes the operation of such devices with the operation of the system as a whole.
- 2) For zSeries, a storage server with ESCON or OEMI interfaces. The control unit adapts a native device interface to an I/O interface that a zSeries host system supports.
- 3) The portion of the storage unit that supports the attachment of emulated count key data devices over ESCON, FICON, or OEMI interfaces. See also *cluster*.

control-unit image

For zSeries, a logical subsystem that is accessed through an ESCON I/O interface. One or more control-unit images exist in each control unit. Each image appears as an independent control unit, but all control-unit images share a common set of hardware facilities. The

DS8000 can emulate 3990-3, TPF, 3990-6, or 2105 control units.

control-unit-initiated reconfiguration (CUIR)

A software mechanism that the DS8000 uses to request that an operating system of a zSeries host verify that one or more subsystem resources can be taken offline for service. The DS8000 can use this process to automatically vary channel paths offline and online to facilitate bay service or concurrent code installation. Depending on the operating system, support for this process might be model dependent, might depend on the IBM TotalStorage Enterprise Storage Server Subsystem Device Driver, or might not exist.

Coordinated Universal Time (UTC)

The international standard of time that is kept by atomic clocks around the world.

Copy Services

A collection of optional software features, with a Web-browser interface, used for configuring, managing, and monitoring data-copy functions.

Copy Services CLI

See *Copy Services command-line interface*.

Copy Services client

Software that runs on each DS8000 cluster in the Copy Services server group and that performs the following functions:

- Communicates configuration, status and connectivity information to the Copy Services server
- Performs data-copy functions on behalf of the Copy Services server

Copy Services command-line interface (Copy Services CLI)

The command-line interface software that is provided with DS8000 Copy Services and used for invoking Copy Services functions from host systems attached to the DS8000. See also *command-line interface*.

Copy Services server

A cluster that the Copy Services administrator designates to perform the DS8000 Copy Services functions.

Copy Services server group

A collection of user-designated DS8000

clusters participating in Copy Services functions that a designated, active, Copy Services server manages. A Copy Services server group is also called a Copy Services domain.

count field

The first field of a count key data (CKD) record. This eight-byte field contains a four-byte track address (CCHH). It defines the cylinder and head that are associated with the track, and a one-byte record number (R) that identifies the record on the track. It defines a one-byte key length that specifies the length of the record's key field (0 means no key field). It defines a two-byte data length that specifies the length of the record's data field (0 means no data field). Only the end-of-file record has a data length of zero.

count key data (CKD)

For zSeries, a data-record format employing self-defining record formats in which each record is represented by up to three fields: a *count* field that identifies the record and specifies its format, an optional *key* field that identifies the data area contents, and an optional *data* field that typically contains the user data. For CKD records on the storage unit, the logical volume size is defined in terms of the device emulation mode (3390 or 3380 track format). The count field is always 8 bytes long and contains the lengths of the key and data fields, the key field has a length of 0 to 255 bytes, and the data field has a length of 0 to 65 535 or the maximum that will fit on the track. See also *data record*.

CPC See *cluster processor complex*.

CRC See *cyclic redundancy check*.

CU See *control unit*.

CUIR See *control-unit initiated reconfiguration*.

custom volume

A volume in count-key-data (CKD) format that is not a standard volume, which means that it does not necessarily present the same number of cylinders and capacity to its assigned logical control unit as provided by one of the following standard zSeries volume types: 3390-2, 3390-3, 3390-9, 3390-2 (3380-track

mode), or 3390-3 (3380-track mode). See also *count-key-data*, *interleave*, *standard volume*, and *volume*.

CUT See *Coordinated Universal Time*.

cyclic redundancy check (CRC)

A redundancy check in which the check key is generated by a cyclic algorithm. (T)

cylinder

A unit of storage on a CKD device with a fixed number of tracks.

D

DA See *device adapter*. See also *IBM Serial Storage adapter*.

daisy chain

See *serial connection*.

DASD See *direct access storage device*.

DASD fast write (DFW)

A function of a storage server in which active write data is stored in nonvolatile cache, thus avoiding exposure to data loss.

data availability

The degree to which data is available when needed, typically measured as a percentage of time that the system would be capable of responding to any data request (for example, 99.999% available).

data compression

A technique or algorithm used to encode data such that the encoded result can be stored in less space than the original data. The original data can be recovered from the encoded result through a reverse technique or reverse algorithm. See also *compression*.

Data Facility Storage Management Subsystem (DFSMS)

An operating environment that helps automate and centralize the management of storage. To manage storage, DFSMS provides the storage administrator with control over data class, storage class, management class, storage group, and automatic class selection routine definitions.

data field

The optional third field of a count key data (CKD) record. The count field specifies the

length of the data field. The data field contains data that the program writes.

data record

The basic unit of zSeries storage on a DS8000, also known as a count-key-data (CKD) record. Data records are stored on a track. The records are sequentially numbered starting with 0. The first record, R0, is typically called the track descriptor record and contains data that the operating system normally uses to manage the track. See also *count-key-data* and *fixed-block architecture*.

data set FlashCopy

For zSeries hosts, a feature of FlashCopy that indicates how many partial volume FlashCopy relationships are active on a volume.

data sharing

The ability of multiple host systems to concurrently utilize data that they store on one or more storage devices. The storage unit enables configured storage to be accessible to any, or all, attached host systems. To use this capability, the host program must be designed to support data that it is sharing.

DDM See *disk drive module*.

DDM group

See *disk pack*.

dedicated storage

Storage within a storage unit that is configured such that a single host system has exclusive access to the storage.

demote

To remove a logical data unit from cache memory. A storage server demotes a data unit to make room for other logical data units in the cache or because the logical data unit is not valid. The storage unit must destage logical data units with active write units before they can be demoted. See also *destage*.

destage

To move data from an online or higher priority to an offline or lower priority device. The storage unit stages incoming data into cache and then destages it to disk.

device

For zSeries, a disk drive.

device adapter (DA)

A physical component of the DS8000 that provides communication between the clusters and the storage devices. The DS8000 has eight device adapters that it deploys in pairs, one from each cluster. Device adapter pairing enables the DS8000 to access any disk drive from either of two paths, providing fault tolerance and enhanced availability.

device address

For zSeries, the field of an ESCON device-level frame that selects a specific device on a control-unit image.

device ID

The unique two-digit hexadecimal number that identifies the logical device.

device interface card

A physical subunit of a storage cluster that provides the communication with the attached device drive modules.

device number

For zSeries, a four-hexadecimal-character identifier, for example 13A0, that the systems administrator associates with a device to facilitate communication between the program and the host operator. The device number is associated with a subchannel.

device sparing

A subsystem function that automatically copies data from a failing device drive module to a spare device drive module. The subsystem maintains data access during the process.

DFS See *distributed file service*.

DFSMS

See *Data Facility Storage Management Subsystem*.

direct access storage device (DASD)

- 1) A mass storage medium on which a computer stores data.
- 2) A disk device.

disk cage

A container for disk drives. Each disk cage supports eight disk packs (64 disks).

disk drive

Standard term for a disk-based nonvolatile

storage medium. The DS8000 uses hard disk drives as the primary nonvolatile storage media to store host data.

disk drive module (DDM)

A field replaceable unit that consists of a single disk drive and its associated packaging.

disk drive module group

See *disk pack*.

disk pack

A group of disk drive modules (DDMs) installed as a unit in a DDM bay.

disk group

A collection of 4 disk drives that are connected to the same pair of IBM Serial Storage adapters and can be used to create a RAID array. A disk group can be formatted as count key data or fixed block, and as RAID or non-RAID, or it can be left unformatted. A disk group is a logical assemblage of disk drives. Contrast with *disk pack*.

distributed file service (DFS)

A service that provides data access over IP networks.

DNS See *domain name system*.

domain

- 1) That part of a computer network in which the data processing resources are under common control.
- 2) In TCP/IP, the naming system used in hierarchical networks.
- 3) A Copy Services server group, in other words, the set of clusters the user designates to be managed by a particular Copy Services server.

domain name system (DNS)

In TCP/IP, the server program that supplies name-to-address translation by mapping domain names to internet addresses. The address of a DNS server is the internet address of the server that hosts the DNS software for the network.

dotted decimal notation

A convention used to identify IP addresses. The notation consists of four 8-bit numbers written in base 10. For example, 9.113.76.250 is an IP address that contains the octets 9, 113, 76, and 250.

drawer

A unit that contains multiple device drive modules and provides power, cooling, and related interconnection logic to make the device drive modules accessible to attached host systems.

drive 1) A peripheral device, especially one that has addressed storage media. See also *disk drive module*.

2) The mechanism used to seek, read, and write information on a storage medium.

DS8000

See *IBM TotalStorage DS8000*.

DS8000 Batch Configuration tool

A program that automatically configures a DS8000. The configuration is based on data that IBM service personnel enter into the program.

DS Storage Manager

See *IBM TotalStorage DS Storage Manager*.

duplex

- 1) Regarding Copy Services, the state of a volume pair after Remote Mirror and Copy has completed the copy operation and the volume pair is synchronized.
- 2) In general, pertaining to a communication mode in which data can be sent and received at the same time.

dynamic sparing

The ability of a storage server to move data from a failing disk drive module (DDM) to a spare DDM while maintaining storage functions.

E

E10 The predecessor of the F10 model of the Enterprise Storage Server. See also *F10*.

E20 The predecessor of the F20 model of the Enterprise Storage Server. See also *F20*.

EBCDIC

See *extended binary-coded decimal interchange code*.

EC See *engineering change*.

ECKD See *extended count key data*.

eight pack

See *disk pack*.

electrostatic discharge (ESD)

An undesirable discharge of static electricity that can damage equipment and degrade electrical circuitry.

emergency power off (EPO)

A means of turning off power during an emergency, usually a switch.

EMIF See *ESCON multiple image facility*.

enclosure

A unit that houses the components of a storage subsystem, such as a control unit, disk drives, and power source.

end of file

A coded character recorded on a data medium to indicate the end of the medium. On a count-key-data direct access storage device, the subsystem indicates the end of a file by including a record with a data length of zero.

engineering change (EC)

An update to a machine, part, or program.

Enterprise Systems Architecture/390 (ESA/390)

An IBM architecture for mainframe computers and peripherals. Processor systems that follow the ESA/390 architecture include the ES/9000® family. See also *z/Architecture*.

Enterprise Systems Connection (ESCON)

1) A zSeries computer peripheral interface. The I/O interface uses zSeries logical protocols over a serial interface that configures attached units to a communication fabric.

2) A set of IBM products and services that provide a dynamically connected environment within an enterprise.

EPO See *emergency power off*.

ERDS See *error-recording data set*.

ERP See *error recovery procedure*.

error-recording data set (ERDS)

On zSeries hosts, a data set that records data-storage and data-retrieval errors. A service information message (SIM) provides the error information for the ERDS.

error recovery procedure (ERP)

Procedures designed to help isolate and, where possible, to recover from errors in equipment. The procedures are often used

in conjunction with programs that record information on machine malfunctions.

ESA/390

See *Enterprise Systems Architecture/390*.

ESCD See *ESCON director*.

ESCON

See *Enterprise System Connection*.

ESCON channel

A zSeries channel that supports ESCON protocols.

ESCON director (ESCD)

An I/O interface switch that allows the interconnection of multiple ESCON interfaces in a distributed-star topology.

ESCON host systems

zSeries hosts that attach to the DS8000 with an ESCON adapter. Such host systems run on operating systems that include MVS, VSE, TPF, or versions of VM.

ESCON multiple image facility (EMIF)

For zSeries, a function that enables LPARs to share an ESCON channel path by providing each LPAR with its own channel-subsystem image.

EsconNet

In the DS Storage Manager, the label on a pseudo host icon that represents a host connection that uses the ESCON protocol and that is not completely defined on the DS8000. See also *pseudo host* and *access-any mode*.

ESD See *electrostatic discharge*.

eserver

See *IBM e(logo)server*.

ESSNet

See *IBM TotalStorage Enterprise Storage Server Network*.

extended binary-coded decimal interchange code (EBCDIC)

An IBM-developed coding scheme used to represent various alphabetic, numeric, and special symbols with a coded character set of 256 eight-bit codes.

extended count key data (ECKD)

An extension of the count key data (CKD) architecture.

extent A continuous space on a disk that is occupied by or reserved for a particular

data set, data space, or file. The unit of increment is a track. See also *multiple allegiance* and *parallel access volumes*.

extent pool

A groups of extents. See also *extent*.

F

fabric In fibre-channel technology, a routing structure, such as a switch, receives addressed information and routes to the appropriate destination. A fabric can consist of more than one switch. When multiple fibre-channel switches are interconnected, they are said to be *cascaded*.

failback

Pertaining to a cluster recovery from failover following repair. See also *failover*.

failover

Pertaining to the process of transferring all control to a single cluster when the other cluster in the storage unit fails. See also *cluster* and *failback*.

fast write

A write operation at cache speed that does not require immediate transfer of data to a disk drive. The subsystem writes the data directly to cache, to nonvolatile storage, or to both. The data is then available for destaging. A fast-write operation reduces the time an application must wait for the I/O operation to complete.

FBA See *fixed-block architecture*.

FC See *feature code*. **Note:** *FC* is a common abbreviation for fibre channel in the industry, but the DS8000 customer documentation library reserves *FC* for feature code.

FC-AL See *Fibre Channel-Arbitrated Loop*.

FCP See *fibre-channel protocol*.

FCS See *fibre-channel standard*.

feature code (FC)

A code that identifies a particular orderable option and that is used by service personnel to process hardware and software orders. Individual optional features are each identified by a unique feature code.

fibre channel

A data-transmission architecture based on the ANSI Fibre Channel standard, which supports full-duplex communication. The DS8000 supports data transmission over fiber-optic cable through its fibre-channel adapters. See also *fibre-channel protocol* and *fibre-channel standard*.

Fibre Channel-Arbitrated Loop (FC-AL)

An implementation of the Fibre Channel standard that uses a ring topology for the communication fabric. Refer to American National Standards Institute (ANSI) X3T11/93-275. In this topology, two or more fibre-channel end points are interconnected through a looped interface.

fibre-channel connection (FICON)

A fibre-channel communications protocol designed for IBM mainframe computers and peripherals.

fibre-channel protocol (FCP)

A protocol used in fibre-channel communications with five layers that define how fibre-channel ports interact through their physical links to communicate with other ports.

Fibre-Channel standard (FCS)

An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric. The protocol has two layers. The IP layer defines basic interconnection protocols. The upper layer supports one or more logical protocols (for example, FCP for SCSI command protocols and SBICON for zSeries command protocols). Refer to American National Standards Institute (ANSI) X3.230-199x. See also *fibre-channel protocol*.

fibre-channel topology

An interconnection topology supported on fibre-channel adapters. See also *point-to-point connection*, *switched fabric*, and *arbitrated loop*.

FICON

See *fibre-channel connection*.

FiconNet

In the DS Storage Manager, the label on a pseudo host icon that represents a host connection that uses the FICON protocol

and that is not completely defined on the DS8000. See also *pseudo host* and *access-any mode*.

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. (GC)

FIFO See *first-in-first-out*.

File Transfer Protocol (FTP)

In TCP/IP, an application protocol used to transfer files to and from host computers. See also *Transmission Control Protocol/Internet Protocol*.

firewall

A protection against unauthorized connection to a computer or a data storage system. The protection is usually in the form of software on a gateway server that grants access to users who meet authorization criteria.

first-in-first-out (FIFO)

A queuing technique in which the next item to be retrieved is the item that has been in the queue for the longest time. (A)

fixed-block architecture (FBA)

An architecture for logical devices that specifies the format of and access mechanisms for the logical data units on the device. The logical data unit is a block. All blocks on the device are the same size (fixed size). The subsystem can access them independently.

fixed-block device

An architecture for logical devices that specifies the format of the logical data units on the device. The logical data unit is a block. All blocks on the device are the same size (fixed size); the subsystem can access them independently. This is the required format of the logical data units for host systems that attach with a SCSI or fibre-channel interface. See also *fibre channel* and *small computer systems interface*.

FlashCopy

An optional feature of the DS8000 that can make an instant copy of data, that is, a point-in-time copy of a volume.

FlashCopy relationship

A mapping of a FlashCopy source volume

and a FlashCopy target volume that allows a point-in-time copy of the source volume to be copied to the target volume. FlashCopy relationships exist from the time that you initiate a FlashCopy operation until the storage unit copies all data from the source volume to the target volume or until you delete the FlashCopy relationship, if it is persistent.

FRU See *field replaceable unit*.

FTP See *File Transfer Protocol*.

full duplex

See *duplex*.

fuzzy copy

A function of the Global Copy feature wherein modifications to the primary logical volume are performed on the secondary logical volume at a later time. The original order of update is not strictly maintained. See also *Global Copy*.

G

GB See *gigabyte*.

GDPS See *Geographically Dispersed Parallel Sysplex*.

Geographically Dispersed Parallel Sysplex (GDPS)

A zSeries multisite application-availability solution.

gigabyte (GB)

A gigabyte of storage is 10^9 bytes. A gigabyte of memory is 2^{30} bytes.

Global Copy

An optional capability of the DS8000 remote mirror and copy feature that maintains a fuzzy copy of a logical volume on the same DS8000 or on another DS8000. In other words, all modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume at a later point in time. The original order of update is not strictly maintained. See also *Remote Mirror and Copy* and *Metro Mirror*.

Global Mirror

An optional capability of the DS8000 remote mirror and copy feature that provides a 2-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 *Metro Mirror* and *Remote Mirror and Copy*.

group In DS8000 documentation, a nickname for two different kinds of groups, depending on the context. See *disk pack* or *Copy Services server group*.

H

HA See *host adapter*.

HACMP

See *High-Availability Cluster Multi-Processing*.

hard disk drive (HDD)

- 1) A storage medium within a storage server used to maintain information that the storage server requires.
- 2) A mass storage medium for computers that is typically available as a fixed disk (such as the disks used in system units of personal computers or in drives that are external to a personal computer) or a removable cartridge.

hardware service manager (HSM)

An option on an AS/400 or iSeries host that enables the user to display and work with system hardware resources and to debug input-output processors (IOP), input-output adapters (IOA), and devices.

HCD See *Hardware Configuration Data*.

HDA See *head and disk assembly*.

HDD See *hard disk drive*.

hdisk An AIX term for storage space.

head and disk assembly (HDA)

The portion of an HDD associated with the medium and the read/write head.

heartbeat

A status report sent at regular intervals from the DS8000. The service provider uses this report to monitor the health of the call home process. See also *call home*, *heartbeat call home record*, and *remote technical assistance information network*.

heartbeat call home record

Machine operating and service information sent to a service machine. These records might include such information as feature

code information and product logical configuration information.

hierarchical storage management

1) A function in storage management software, such as Tivoli Storage Management or Data Facility Storage Management Subsystem/MVS (DFSMS/MVS), that automatically manages free space based on the policy that the storage administrator sets.

2) In AS/400 storage management, an automatic method to manage and distribute data between the different storage layers, such as disk units and tape library devices.

High-Availability Cluster Multi-Processing (HACMP)

Software that provides host clustering, so that a failure of one host is recovered by moving jobs to other hosts within the cluster.

high-speed link (HSL)

A hardware connectivity architecture that links system processors to system input/output buses and other system units.

home address

A nine-byte field at the beginning of a track that contains information that identifies the physical track and its association with a cylinder.

hop Interswitch connection. A hop count is the number of connections that a particular block of data traverses between source and destination. For example, data traveling from one hub over a wire to another hub traverses one hop.

host See *host system*.

host adapter

A physical subunit of a storage server that provides the ability to attach to one or more host I/O interfaces.

host name

The Internet address of a machine in the network. The host name can be entered in the host definition as the fully qualified domain name of the attached host system, such as *mycomputer.city.company.com*, or as the subname of the fully qualified domain name, for example, *mycomputer*. See also *host system*.

host processor

A processor that controls all or part of a user application network. In a network, the processing unit in which the data communication access method resides. See also *host system*.

host system

A computer, either of the mainframe (for example, zSeries) or of the open-systems type, that is connected to the DS8000. zSeries hosts are connected to the DS8000 through ESCON interfaces. Open-systems hosts are connected to the DS8000 by SCSI or fibre-channel interfaces.

hot plug

Pertaining to the ability to add or remove a hardware facility or resource to a unit while power is on.

HSL See *high-speed link*.

HSM See *hierarchical storage management* or *Hardware Service Manager*.

I

i5/OS The IBM operating system that runs the IBM i5/OS and e(logo)Server i5 server families of servers.

IBM e(logo)server

The IBM brand name for a series of server products that are optimized for e-commerce. The products include the iSeries, pSeries, xSeries, and zSeries.

IBM product engineering (PE)

The third-level of IBM service support. Product engineering is composed of IBM engineers who have experience in supporting a product or who are knowledgeable about the product.

IBM Serial Storage adapter

A physical adapter based on the IBM Serial Storage architecture. IBM Serial Storage adapters connect disk drive modules to DS8000 clusters.

IBM TotalStorage

The brand name used to identify storage products from IBM, including the IBM TotalStorage DS8000. See also *IBM TotalStorage DS8000* and *IBM TotalStorage DS Storage Manager*.

IBM TotalStorage DS8000

A member of the IBM TotalStorage

Resiliency Family of storage servers and attached storage devices (disk drive modules). The DS8000 delivers high-performance, fault-tolerant storage and management of enterprise data, affording access through multiple concurrent operating systems and communication protocols. High performance is provided by multiple symmetrical multiprocessors, integrated caching, RAID support for the disk drive modules, and disk access through a high-speed serial storage architecture interface.

IBM TotalStorage DS CLI

The command-line interface (CLI) that is specific to the DS8000.

IBM TotalStorage DS Storage Manager (DS Storage Manager)

Software with a Web-browser interface for configuring the DS8000.

IBM TotalStorage Enterprise Storage Server Network (ESSNet)

A private network providing Web browser access to the Enterprise Storage Server. IBM installs the ESSNet software on an IBM workstation called the IBM TotalStorage ESS Master Console, supplied with the first ESS delivery.

IBM TotalStorage Management Console (MC)

An IBM workstation that acts as the focal point for configuration, Copy Services management, and maintenance for the DS8000. It includes a Web browser that provides links to the user interface, including the DS Storage Manager and the DS8000 Copy Services.

IBM TotalStorage Multipath Subsystem Device Driver (SDD)

Software that is designed to support the multipath configuration environments of the DS8000. The SDD resides in a host system with the native disk device driver.

IBM TotalStorage Resiliency Family

A set of hardware and software features and products, as well as integrated software and services that are available on the IBM TotalStorage DS8000 and the IBM TotalStorage Enterprise Storage Server, Models 750 and 800.

image See *storage image*.

IML See *initial microcode load*.

implicit allegiance

In Enterprise Systems Architecture/390, a relationship that a control unit creates between a device and a channel path when the device accepts a read or write operation. The control unit guarantees access to the channel program over the set of channel paths that it associates with the allegiance.

initial microcode load (IML)

The action of loading microcode for a computer into that computer's storage.

initial program load (IPL)

The action of loading software into a computer, typically an operating system that controls the computer.

initiator

A SCSI device that communicates with and controls one or more targets. Contrast with *target*.

i-node The internal structure in an AIX operating system that describes the individual files in the operating system. It contains the code, type, location, and owner of a file.

input/output (I/O)

Pertaining to (a) input, output, or both or (b) a device, process, or channel involved in data input, data output, or both.

input/output configuration data set

A configuration definition built by the I/O configuration program (IOCP) and stored on disk files associated with the processor controller.

interleave

To automatically create two striped partitions across the drives in a RAID-5 array, both of which use the count-key-data (CKD) record format.

Internet Protocol (IP)

In the Internet suite of protocols, a protocol without connections that routes data through a network or interconnecting networks and acts as an intermediary between the higher protocol layers and the physical network. The upper layer supports one or more logical protocols (for example, a SCSI-command protocol and a zSeries command protocol). Refer to ANSI X3.230-199x. The IP acronym is the IP in TCP/IP. See also *Transmission Control Protocol/Internet Protocol*.

invalidate

To remove a logical data unit from cache memory because it cannot support continued access to the logical data unit on the device. This removal might be the result of a failure within the storage server or a storage device that is associated with the device.

I/O See *input/output*.

I/O adapter (IOA)

An input-output adapter on the PCI bus.

IOCDs

See *input/output configuration data set*.

IOCP See *I/O Configuration Program*.

I/O Configuration Program (IOCP)

A program that defines to a system all the available I/O devices and channel paths.

I/O device

An addressable read and write unit, such as a disk drive device, magnetic tape device, or printer.

I/O interface

An interface that enables a host to perform read and write operations with its associated peripheral devices.

I/O Priority Queueing

A facility in the Workload Manager of zSeries that enables the system administrator to set priorities for queueing I/Os from different system images. See also *multiple allegiance* and *parallel access volume*.

I/O processor (IOP)

Controls input-output adapters and other devices.

I/O sequential response time

The time an I/O request is queued in processor memory waiting for previous I/Os to the same volume to complete.

IOSQ See *I/O sequential response time*.

IP See *Internet Protocol*.

IPL See *initial program load*.

iSeries

An IBM e(logo)server product that emphasizes integration. It is the successor to the AS/400 family of servers.

J

Java Virtual Machine (JVM)

A software implementation of a central processing unit (CPU) that runs compiled Java code (applets and applications).
(GC)

JVM See *Java Virtual Machine*.

K

KB See *kilobyte*.

key field

The second (optional) field of a count key data record. The key length is specified in the count field. The key length determines the field length. The program writes the data in the key field and uses the key field to identify or locate a given record. The subsystem does not use the key field.

kilobyte (KB)

1) For processor storage, real, and virtual storage, and channel volume, 2^{10} or 1024 bytes.

2) For disk storage capacity and communications volume, 1000 bytes.

Korn shell

Interactive command interpreter and a command programming language.

KPOH See *thousands of power-on hours*.

L

LAN See *local area network*.

last-in first-out (LIFO)

A queuing technique in which the next item to be retrieved is the item most recently placed in the queue. (A)

LBA See *logical block address*.

LCU See *logical control unit*.

least recently used (LRU)

1) The algorithm used to identify and make available the cache space that contains the least-recently used data.

2) A policy for a caching algorithm that chooses to remove from cache the item that has the longest elapsed time since its last access.

LED See *light-emitting diode*.

LMC See *licensed machine code*.

licensed machine code (LMC)

Microcode that IBM does not sell as part of a machine, but licenses to the customer. LMC is implemented in a part of storage that is not addressable by user programs. Some IBM products use it to implement functions as an alternate to hard-wired circuitry.

LIFO See *last-in first-out*.

light-emitting diode (LED)

A semiconductor chip that gives off visible or infrared light when activated.

link address

On an ESCON interface, the portion of a source or destination address in a frame that ESCON uses to route a frame through an ESCON director. ESCON associates the link address with a specific switch port that is on the ESCON director. Equivalently, it associates the link address with the channel subsystem or control unit link-level functions that are attached to the switch port.

link-level facility

The ESCON hardware and logical functions of a control unit or channel subsystem that allow communication over an ESCON write interface and an ESCON read interface.

local area network (LAN)

A computer network located on a user's premises within a limited geographic area.

local e-mail

An e-mail configuration option for storage servers that are connected to a host-system network that does not have a domain name system (DNS) server.

logical address

On an ESCON interface, the portion of a source or destination address in a frame used to select a specific channel-subsystem or control-unit image.

logical block address (LBA)

The address assigned by the DS8000 to a sector of a disk.

logical control unit (LCU)

See *control-unit image*.

logical data unit

A unit of storage that is accessible on a given device.

logical device

The facilities of a storage server (such as the DS8000) associated with the processing of I/O operations directed to a single host-accessible emulated I/O device. The associated storage is referred to as a logical volume. The logical device is mapped to one or more host-addressable units, such as a device on a zSeries I/O interface or a logical unit on a SCSI I/O interface, such that the host initiating I/O operations to the I/O-addressable unit interacts with the storage on the associated logical device.

logical partition (LPAR)

For zSeries, a set of functions that create the programming environment in which more than one logical partition (LPAR) is established on a processor. An LPAR is conceptually similar to a virtual machine environment except that the LPAR is a function of the processor. Also, the LPAR does not depend on an operating system to create the virtual machine environment.

logical path

1) The relationship between a channel image and a control-unit image that designates the physical path to be used for device-level communications between these images. The logical path is established as part of the channel and control-unit initialization procedures by the exchange of link-level frames.

2) With the Remote Mirror and Copy feature, the relationship between a source logical subsystem (LSS) and a target LSS that is created over a physical path through the interconnection fabric that is used for Remote Mirror and Copy functions. An LSS is a primary control unit, which performs the functions of a channel image.

logical subsystem (LSS)

A topological construct that consists of a group of up to 256 logical devices. A DS8000 can have up to 16 CKD-formatted logical subsystems (4096 CKD logical devices) and also up to 16 fixed-block logical subsystems (4096 fixed-block logical devices). The logical subsystem facilitates configuration of the DS8000 and might have other implications relative to the operation of certain functions. There is

a one-to-one mapping between a CKD logical subsystem and a zSeries control-unit image.

For zSeries hosts, a logical subsystem represents a logical control unit (LCU). Each control-unit image is associated with only one logical subsystem. See also *control-unit image*.

logical unit

In open systems, a logical disk drive.

logical unit number (LUN)

In the SCSI protocol, a unique number that is used on a SCSI bus to enable it to differentiate between separate devices, each of which is a logical unit.

logical volume

The storage medium that is associated with a logical disk drive. A logical volume typically resides on one or more storage devices. The DS8000 administrator defines this unit of storage. The logical volume, when residing on a RAID-formatted array, is spread over the drives in the array.

logical volume manager (LVM)

A set of system commands, library routines, and other tools that allow the user to establish and control logical volume storage. The LVM maps data between the logical view of storage space and the physical disk drive module.

longitudinal redundancy check (LRC)

1) A method of error checking during data transfer that involves checking parity on a row of binary digits that are members of a set that forms a matrix. Longitudinal redundancy check is also called a longitudinal parity check.

2) A mechanism that the DS8000 uses for locating errors. The LRC checks the data as it progresses from the host, through the DS8000 controller, into the device adapter, and to the array.

longwave laser adapter

A connector that is used between a host and the DS8000 to support longwave fibre-channel communication.

loop

The physical connection between a pair of device adapters in the DS8000. See also *device adapter*.

LPAR See *logical partition*.

LRC See *longitudinal redundancy check*.

LRU See *least recently used*.

LSS See *logical subsystem*.

LUN See *logical unit number*.

LVM See *logical volume manager*.

M

machine level control (MLC)

A database that contains the EC level and configuration of products in the field.

machine reported product data (MRPD)

Product data gathered by a machine and sent to a destination such as an IBM support server or RETAIN. These records might include such information as feature code information and product logical configuration information.

mainframe

A computer, usually in a computer center, with extensive capabilities and resources to which other computers may be connected so that they can share facilities. (T)

maintenance analysis procedure (MAP)

A hardware maintenance document that gives an IBM service representative a step-by-step procedure for tracing a symptom to the cause of a failure.

management console

See *IBM TotalStorage Management Console*.

Management Information Base (MIB)

1) A collection of objects that can be accessed by means of a network management protocol. (GC)

2) The MIB record conforms to the Open Systems Interconnection (OSI) standard defined by the International Organization for Standardization (ISO) for the exchange of information. See also *simple network management protocol*.

MAP See *maintenance analysis procedure*.

master storage unit

The physical unit that controls the creation of consistency groups in a Global Mirror session. The master storage unit sends commands to subordinate storage units. A storage unit can be a master for only one

Global Mirror session. Contrast with *subordinate storage unit*.

maximum consistency group drain time

The value in seconds that indicates the maximum time that writes from the local site are delayed to the remote site while the current consistency group is being formed at the remote site. When this time is exceeded, the current attempt to form a consistency group is ended and another attempt is started. If this time is exceeded five times, this maximum time is ignored on the next attempt to form a consistency group. The default value is the larger of four minutes or two times the consistency group interval time if this value is set to zero.

maximum coordination time

The value in milliseconds that indicates the maximum time that is allowed for host I/O to be delayed during the coordination of the primary volumes of an Global Mirror session. The default is 50 milliseconds if this value is set to zero.

MB See *megabyte*.

MC See *IBM TotalStorage Management Console*.

MCA See *Micro Channel architecture*.

MDM See *Multiple Device Manager*.

mean time between failures (MTBF)

1) A projection of the time that an individual unit remains functional. The time is based on averaging the performance, or projected performance, of a population of statistically independent units. The units operate under a set of conditions or assumptions.

2) For a stated period in the life of a functional unit, the mean value of the lengths of time between consecutive failures under stated conditions. (I) (A)

medium

For a storage unit, the disk surface on which data is stored.

megabyte (MB)

1) For processor storage, real and virtual storage, and channel volume, 2^{20} or 1 048 576 bytes.

2) For disk storage capacity and communications volume, 1 000 000 bytes.

Metro Mirror

A function of a storage server that maintains a consistent copy of a logical volume on the same storage server or on another storage server. All modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume. See also *Remote Mirror and Copy* and *Global Copy*.

MES See *miscellaneous equipment specification*.

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. The architecture defines the services that each subsystem can or must provide.

Microsoft Internet Explorer (MSIE)

Web browser software manufactured by Microsoft.

migration

The replacement of a system or subsystem with a different type of system or subsystem, such as replacing a SCSI host adapter with a fibre-channel host adapter. In the context of data migration regarding the DS8000, the transfer of data from one storage unit to another, such as from a 3390 to the DS8000.

MIH See *missing-interrupt handler*.

mirrored pair

Two units that contain the same data. The system refers to them as one entity.

mirroring

In host systems, the process of writing the same data to two disk units within the same auxiliary storage pool at the same time.

miscellaneous equipment specification (MES)

IBM field-installed change to a machine.

missing-interrupt handler (MIH)

An MVS and MVS/XA facility that tracks I/O interrupts. MIH informs the operator and creates a record whenever an expected interrupt fails to occur before a specified elapsed time is exceeded.

MLC See *machine level control*.

mobile solutions terminal (MoST)

The mobile terminal used by service personnel.

mode conditioning patch cable

A cable that converts a single-mode signal from a longwave adapter into a light signal that is appropriate for multimode fibre. Another mode conditioning patch cable is required at the terminating end of the multimode fibre to convert the signal back to a single-mode signal for a longwave adapter.

Model 100

A 2105 Model 100, often simply referred to as a Mod 100, is an expansion enclosure for the Enterprise Storage Server. See also *2105*.

MoST See *mobile solutions terminal*.

MRPD See *machine reported product data*.

MSA See *multiport serial adapter*.

MSIE See *Microsoft Internet Explorer*.

MTBF See *mean time between failures*.

Multipath Subsystem Device Driver

See *IBM TotalStorage DS8000 Multipath Subsystem Device Driver*.

multiple allegiance

A DS8000 hardware function that is independent of software support. This function enables multiple system images to concurrently access the same logical volume on the DS8000 as long as the system images are accessing different extents. See also *extent* and *parallel access volumes*.

Multiple Device Manager (MDM)

A component of the IBM TotalStorage Productivity Center that allows administrators to configure, manage, and monitor the performance of SAN storage devices from a single console.

multiple relationship FlashCopy

An option of the DS8000 that creates backup copies from one source to multiple targets by simultaneously establishing multiple FlashCopy relationships.

multiple virtual storage (MVS)

Implies MVS/390, MVS/XA, MVS/ESA, and the MVS element of the zSeries operating system.

multiplex

The action of transmitting simultaneously.

multiport serial adapter (MSA)

An adapter on the IBM TotalStorage Management Console that has multiple ports to which aDS8000 can be attached.

multiprocessor

A computer that includes two or more processors that have common access to a main storage. For the DS8000, the multiprocessors operate in parallel.

MVS See *multiple virtual storage*.

N**name server**

A server that stores names of the participating DS8000 clusters.

Netfinity

IBM Intel-processor-based server; predecessor to the IBM xSeries server.

Netscape Navigator

Web browser software manufactured by Netscape.

network manager

A program or group of programs that is used to monitor, manage, and diagnose the problems of a network. (GC)

node The unit that is connected in a fibre-channel network. A DS8000 is a node in a fibre-channel network.

non-RAID

A disk drive set up independently of other disk drives and not set up as part of a disk pack to store data using the redundant array of disks (RAID) data-striping methodology.

nonremovable medium

A recording medium that cannot be added to or removed from a storage device.

nonvolatile storage (NVS)

Memory that stores active write data to avoid data loss in the event of a power loss.

NVS See *nonvolatile storage*.

O

octet In Internet Protocol addressing, one of the

four parts of a 32-bit integer presented in dotted decimal notation. See also *dotted decimal notation*.

OEMI See *original equipment manufacturer's information*.

open system

A system whose characteristics comply with standards made available throughout the industry and that therefore can be connected to other systems complying with the same standards. Applied to the DS8000, such systems are those hosts that connect to the DS8000 through SCSI or FCP protocols. See also *small computer system interface* and *fibre-channel protocol*.

organizationally unique identifier (OUI)

An IEEE-standards number that identifies an organization with a 24-bit globally unique assigned number referenced by various standards. OUI is used in the family of 802 LAN standards, such as Ethernet and Token Ring.

original equipment manufacturer's information (OEMI)

A reference to an IBM guideline for a computer peripheral interface. The interface uses ESA/390 logical protocols over an I/O interface that configures attached units in a multidrop bus topology.

OS/390

The IBM operating system that includes and integrates functions that many IBM software products (including the MVS operating system) previously provided for the IBM S/390 family of enterprise servers.

OUI See *organizationally unique identifier*.

P

panel The formatted display of information that appears on a display screen.

parallel access volume (PAV)

An advanced function of the DS8000 that enables OS/390 and z/OS systems to issue concurrent I/O requests against a count key data logical volume by associating multiple devices of a single control-unit image with a single logical device. Up to eight device addresses can be assigned to a PAV. The PAV function enables two or more concurrent write

operations to the same logical volume, as long as the write operations are not to the same extents. See also *extent*, *I/O Priority Queueing*, and *multiple allegiance*.

parity A data checking scheme used in a computer system to ensure the integrity of the data. The RAID implementation uses parity to re-create data if a disk drive fails.

path group

In zSeries architecture, a set of channel paths that are defined to a control unit as being associated with a single logical partition (LPAR). The channel paths are in a group state and are online to the host. See also *logical partition*.

path group identifier

In zSeries architecture, the identifier that uniquely identifies a given logical partition (LPAR). The path group identifier is used in communication between the LPAR program and a device. The identifier associates the path group with one or more channel paths, thereby defining these paths to the control unit as being associated with the same LPAR. See also *logical partition*.

PAV See *parallel access volume*.

PCI See *peripheral component interconnect*.

PDU See *protocol data unit*.

PE See *IBM product engineering*.

peripheral component interconnect (PCI)

An architecture for a system bus and associated protocols that supports attachments of adapter cards to a system backplane.

persistent FlashCopy

A state where a FlashCopy relationship remains indefinitely until the user deletes it. The relationship between the source and target volumes is maintained after a background copy completes.

physical path

A single path through the I/O interconnection fabric that attaches two units. For Copy Services, this is the path from a host adapter on one DS8000 (through cabling and switches) to a host adapter on another DS8000.

pinned data

Data that is held in cache until either an

error condition is corrected and it can be moved to disk storage or until the data is discarded by a host command. Pinned data conditions can only occur on an ESS Model 800 during fast-write or dual-copy functions.

point-in-time copy

A FlashCopy option that creates an instantaneous view of original source data at a specific moment in time.

point-to-point connection

A fibre-channel topology that enables the direct interconnection of ports. See also *arbitrated loop* and *switched fabric*.

port

A physical connection on a host adapter to the cable that connects the DS8000 to hosts, switches, or another DS8000. The DS8000 uses SCSI and ESCON host adapters that have two ports per adapter, and fibre-channel host adapters that have one port. See also *ESCON*, *fibre channel*, *host adapter*, and *small computer system interface*.

POST See *power-on self test*.

power-on self test (POST)

A diagnostic test that servers or computers run when they are turned on.

predictable write

A write operation that can cache without knowledge of the existing format on the medium. All write operations on FBA DASD devices are predictable. On CKD DASD devices, a write operation is predictable if it does a format write operation for the first data record on the track.

primary control unit

The DS8000 to which a Remote Mirror and Copy primary device is physically attached.

processor complex

A partition of a storage server that is capable of performing all defined functions of the storage server. Multiple processor complexes provide redundancy.

product engineering

See *IBM product engineering*.

program

On a computer, a generic term for software that controls the operation of the computer. Typically, the program is a

logical assemblage of software modules that perform multiple related tasks.

program-controlled interruption

An interruption that occurs when an I/O channel fetches a channel command word with the program-controlled interruption flag on.

program temporary fix (PTF)

A temporary solution to, or bypass of, a problem diagnosed by IBM as the result of a defect in a current unaltered release of a licensed program. (GC)

promote

To add a logical data unit to cache memory.

protected volume

In AS/400, a disk storage device that is protected from data loss by RAID techniques. An AS/400 host does not mirror a volume configured as a protected volume, while it does mirror all volumes configured as unprotected volumes. The DS8000, however, can be configured to indicate that an AS/400 volume is protected or unprotected and give it RAID protection in either case.

protocol data unit (PDU)

A unit of data specified in the protocol of a given layer and consisting of protocol control information for the layer and, possibly, user data for the layer.

pSeries

The product name of an IBM e(logo)server product that emphasizes performance. It is the successor to the RS/6000 family of servers.

pseudo host

A host connection that is not explicitly defined to the DS8000 and that has access to at least one volume that is configured on the DS8000. The FiconNet pseudo host icon represents the FICON protocol. The EsconNet pseudo host icon represents the ESCON protocol. The pseudo host icon labelled Anonymous represents hosts connected through the FCP protocol. *Anonymous host* is a commonly used synonym for *pseudo host*. The DS8000 adds a pseudo host icon only when it is set to access-any mode. See also *access-any mode*.

PTF See *program temporary fix*.

PV Links

Short for Physical Volume Links, an alternate pathing solution from Hewlett-Packard that provides for multiple paths to a volume, as well as static load balancing.

R

R0 See *track-descriptor record*.

rack See *enclosure*.

RAID See *redundant array of independent disks*. RAID is also commonly expanded to redundant array of *inexpensive* disks. See also *array*.

RAID 5

A type of RAID that optimizes cost-effective performance while emphasizing use of available capacity through data striping. RAID 5 provides fault tolerance for up to two failed disk drives by distributing parity across all the drives in the array plus one parity disk drive. The DS8000 automatically reserves spare disk drives when it assigns arrays to a device adapter pair (DA pair). See also *device adapter*, *RAID 10*, and *redundant array of independent disks*.

RAID 10

A type of RAID that optimizes high performance while maintaining fault tolerance for up to two failed disk drives by striping volume data across several disk drives and mirroring the first set of disk drives on an identical set. The DS8000 automatically reserves spare disk drives when it assigns arrays to a device adapter pair (DA pair). See also *device adapter*, *RAID 5*, and *redundant array of independent disks*.

random access

A mode of accessing data on a medium in a manner that requires the storage device to access nonconsecutive storage locations on the medium.

rank See *array*.

redundant array of independent disks (RAID)

A methodology of grouping disk drives for managing disk storage to insulate data from a failing disk drive.

refresh FlashCopy target volume

An option (previously called *incremental*

FlashCopy) of the DS8000 that creates a point-in-time data copy without copying an entire volume for each point-in-time copy.

Remote Mirror and Copy

A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also *Global Mirror*, *Metro Mirror* and *Global Copy*.

remote technical assistance information network (RETAIN)

The initial service tracking system for IBM service support, which captures heartbeat and call-home records. See also *support catcher* and *support catcher telephone number*.

REQ/ACK

See *request for acknowledgment and acknowledgment*.

request for acknowledgment and acknowledgment (REQ/ACK)

A cycle of communication between two data transport devices for the purpose of verifying the connection, which starts with a request for acknowledgment from one of the devices and ends with an acknowledgment from the second device. The REQ and ACK signals help to provide uniform timing to support synchronous data transfer between an initiator and a target. The objective of a synchronous data transfer method is to minimize the effect of device and cable delays.

reserved allegiance

For zSeries, a relationship that is created in a control unit between a device and a channel path, or path group, when the device completes a Sense Reserve command. The allegiance causes the control unit to guarantee access (that is, busy status is not presented) to the device. Access is over the set of channel paths that are associated with the allegiance; access is for one or more channel programs until the allegiance ends.

RETAIN

See *remote technical assistance information network*.

S

S/390 IBM enterprise servers based on Enterprise Systems Architecture/390 (ESA/390). *S/390* is the currently accepted shortened form of the original name *System/390*.

S/390 storage

Storage arrays and logical volumes that are defined as connected to S/390 servers. This term is synonymous with count-key-data storage.

SAID See *system adapter identification number*.

SAM See *sequential access method*.

SAN See *storage area network*.

SBCON

See *Single-Byte Command Code Sets Connection*.

screen

The physical surface of a display device upon which information is shown to users.

SCSI See *small computer system interface*.

SCSI device

A disk drive connected to a host through an I/O interface using the SCSI protocol. A SCSI device is either an initiator or a target. See also *initiator* and *small computer system interface*.

SCSI-FCP

Synonym for fibre-channel protocol, a protocol used to transport data between an open-systems host and a fibre-channel adapter on an DS8000. See also *fibre-channel protocol* and *small computer system interface*.

SCSI host systems

Host systems that are attached to the DS8000 with a SCSI interface. Such host systems run on UNIX, i5/OS, Windows NT, Windows 2000, or Novell NetWare operating systems.

SCSI ID

A unique identifier assigned to a SCSI device that is used in protocols on the SCSI interface to identify or select the device. The number of data bits on the SCSI bus determines the number of available SCSI IDs. A wide interface has 16 bits, with 16 possible IDs.

SDD See *IBM Subsystem Multipathing Device Driver*.

secondary control unit

The DS8000 to which a Remote Mirror and Copy secondary device is physically attached.

self-timed interface (STI)

An interface that has one or more conductors that transmit information serially between two interconnected units without requiring any clock signals to recover the data. The interface performs clock recovery independently on each serial data stream and uses information in the data stream to determine character boundaries and inter-conductor synchronization.

sequential access

A mode of accessing data on a medium in a manner that requires the storage device to access consecutive storage locations on the medium.

sequential access method (SAM)

An access method for storing, deleting, or retrieving data in a continuous sequence based on the logical order of the records in the file.

serial connection

A method of device interconnection for determining interrupt priority by connecting the interrupt sources serially.

server A host that provides certain services to other hosts that are referred to as clients.

A functional unit that provides services to one or more clients over a network. (GC)

service boundary

A category that identifies a group of components that are unavailable for use when one of the components of the group is being serviced. Service boundaries are provided on the DS8000, for example, in each host bay and in each cluster.

service clearance

The area that is required to open the service covers and to pull out components for servicing.

service information message (SIM)

A message sent by a storage server to service personnel through an zSeries operating system.

service personnel

A generalization referring to individuals or companies authorized to service the DS8000. The terms *service provider*, *service representative*, and *IBM service support representative (SSR)* refer to types of service personnel. See also *service support representative*.

service processor

A dedicated processing unit that is used to service a storage unit.

service support representative (SSR)

Individuals or a company authorized to service the DS8000. This term also refers to a service provider, a service representative, or an IBM service support representative (SSR). An IBM SSR installs the DS8000.

SES SCSI Enclosure Services.

session

A collection of volumes within a logical subsystem that are managed together during the creation of consistent copies of data. All volumes in a session must transfer their data successfully to the remote site before the increment can be called complete.

SFP Small form factor pluggables.

shared storage

Storage that is configured so that multiple hosts can concurrently access the storage. The storage has a uniform appearance to all hosts. The host programs that access the storage must have a common model for the information on a storage device. The programs must be designed to handle the effects of concurrent access.

shortwave laser adapter

A connector that is used between host and DS8000 to support shortwave fibre-channel communication.

SIM See *service information message*.

Simple Network Management Protocol (SNMP)

In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an application layer protocol. Information on devices managed is defined and stored in the application's

Management Information Base (MIB).
(GC) See also *management information base*.

simplex volume

A volume that is not part of a FlashCopy, XRC, or PPRC volume pair.

Single-Byte Command Code Sets Connection (SBCON)

The ANSI standard for the ESCON I/O interface.

small computer system interface (SCSI)

A standard hardware interface that enables a variety of peripheral devices to communicate with one another. (GC)

smart relay host

A mail relay or mail gateway that has the capability to correct e-mail addressing problems.

SMIT See *System Management Interface Tool*.

SMP See *symmetrical multiprocessor*.

SNMP See *Simple Network Management Protocol*.

SNMP agent

A server process that resides on a network node and is responsible for communicating with managers regarding that node. The node is represented as a managed object, which has various fields or variables that are defined in the appropriate MIB.

SNMP manager

A managing system that runs a managing application or suite of applications. These applications depend on Management Information Base (MIB) objects for information that resides on the managed system. Managers generate requests for this MIB information, and an SNMP agent on the managed system responds to these requests. A request can either be the retrieval or modification of MIB information.

software transparency

Criteria applied to a processing environment that states that changes do not require modifications to the host software in order to continue to provide an existing function.

source device

One of the devices in a dual-copy or

remote-copy volume pair. All channel commands to the logical volume are directed to the source device. The data on the source device is duplicated on the target device. See also *target device*.

spare A disk drive on the DS8000 that can replace a failed disk drive. A spare can be predesignated to allow automatic dynamic sparing. Any data preexisting on a disk drive that is invoked as a spare is destroyed by the dynamic sparing copy process.

spatial reuse

A feature of serial storage architecture that enables a device adapter loop to support many simultaneous read/write operations. See also *serial storage architecture*.

SSID See *subsystem identifier*.

SSR See *service support representative*.

stacked status

For zSeries, the condition when the control unit is in a holding status for the channel, and the last time the control unit attempted to present the status, the channel responded with the stack-status control.

stage operation

The operation of reading data from the physical disk drive into the cache.

staging

To move data from an offline or low-priority device back to an online or higher priority device, usually on demand of the system or on request of the user.

standard volume

A volume that emulates one of several zSeries volume types, including 3390-2, 3390-3, 3390-9, 3390-2 (3380-track mode), or 3390-3 (3380-track mode), by presenting the same number of cylinders and capacity to the host as provided by the native zSeries volume type of the same name.

STI See *self-timed interface*.

storage area network

A network that connects a company's heterogeneous storage resources.

storage capacity

The amount of data that a storage

medium can hold; usually expressed in kilobytes, megabytes, or gigabytes.

storage complex

A configuration of one or more storage units that is managed by a management console.

storage device

A physical unit that provides a mechanism to store data on a given medium such that it can be subsequently retrieved. See also *disk drive module*.

storage extent

The minimum contiguous range of storage on a physical storage device, array, or rank that can be allocated to a local volume

storage image

A partitioning of a storage unit that provides emulation of a storage server with one or most storage devices that provides storage capability to a host computer. You can configure more than one storage image on a storage unit.

storage server

A physical unit that manages attached storage devices and provides an interface between them and a host computer by providing the function of one or more logical subsystems. The storage server can provide functions that the storage device does not provide. The storage server has one or more clusters.

storage unit

A physical unit that consists of a storage server that is integrated with one or more storage devices that provide storage capability to a host computer.

storage unit identifier

A unique identifier for a storage unit that consists of a manufacturer, a model number, a type number, a plant of manufacture, and a sequence number.

striping

A technique that distributes data in bit, byte, multibyte, record, or block increments across multiple disk drives.

subagent

An extension to an SNMP agent that permits a user to dynamically add, or in some cases replace, additional management variables in the local MIB,

thereby providing a means of extending the range of information that network managers can access. See also *agent*.

subchannel

A logical function of a channel subsystem associated with the management of a single device.

subordinate storage unit

The physical unit that receives commands from the master storage unit and is specified when a Global Mirror session is started. The subordinate storage unit forms consistency groups and performs other Global Mirror processing. A subordinate storage unit can be controlled by only one master storage unit. Contrast with *master storage unit*.

subsystem identifier (SSID)

A number that uniquely identifies a logical subsystem within a computer installation.

support catcher

See *catcher*.

support catcher telephone number

The telephone number that connects the support catcher server to the DS8000 to receive a trace or dump package. See also *support catcher* and *remote technical assistance information network*.

switched fabric

A fibre-channel topology in which ports are interconnected through a switch. Fabric switches can also be interconnected to support numerous ports on a single network. See also *arbitrated loop* and *point-to-point connection*.

symmetrical multiprocessor (SMP)

An implementation of a multiprocessor computer consisting of several identical processors configured in a way that any subset of the set of processors is capable of continuing the operation of the computer. The DS8000 contains four processors set up in SMP mode.

synchronous write

A write operation whose completion is indicated after the data has been stored on a storage device.

System/390

See *S/390*.

system adapter identification number (SAID)

The unique identification number that is

automatically assigned to each DS8000 host adapter for use by Copy Services.

System Management Interface Tool (SMIT)

An interface tool of the AIX operating system for installing, maintaining, configuring, and diagnosing tasks.

System Modification Program

A program used to install software and software changes on MVS systems.

T

TAP See *Telocator Alphanumeric Protocol*.

target A SCSI device that acts as a subordinate to an initiator and consists of a set of one or more logical units, each with an assigned logical unit number (LUN). The logical units on the target are typically I/O devices. A SCSI target is analogous to a zSeries control unit. See also *small computer system interface*.

target device

One of the devices in a dual-copy or remote-copy volume pair that contains a duplicate of the data that is on the source device. Unlike the source device, the target device might only accept a limited subset of data. See also *source device*.

TB See *terabyte*.

TCP/IP

See *Transmission Control Protocol/Internet Protocol*.

Telocator Alphanumeric Protocol (TAP)

An industry standard protocol for the input of paging requests.

terabyte (TB)

- 1) Nominally, 1 000 000 000 000 bytes, which is accurate when speaking of bandwidth and disk storage capacity.
- 2) For DS8000 cache memory, processor storage, real and virtual storage, a terabyte refers to 2⁴⁰ or 1 099 511 627 776 bytes.

terminal emulator

A function of the management console that allows it to emulate a terminal.

thousands of power-on hours (KPOH)

A unit of time used to measure the mean time between failures (MTBF).

time sharing option (TSO)

An operating system option that provides interactive time sharing from remote terminals.

TotalStorage

See *IBM TotalStorage*.

TPF See *transaction processing facility*.

track A unit of storage on a CKD device that can be formatted to contain a number of data records. See also *home address*, *track-descriptor record*, and *data record*.

track-descriptor record (R0)

A special record on a track that follows the home address. The control program uses it to maintain certain information about the track. The record has a count field with a key length of zero, a data length of 8, and a record number of 0. This record is sometimes referred to as R0.

transaction processing facility (TPF)

A high-availability, high-performance IBM operating system, designed to support real-time, transaction-driven applications. The specialized architecture of TPF is intended to optimize system efficiency, reliability, and responsiveness for data communication and database processing. TPF provides real-time inquiry and updates to a large, centralized database, where message length is relatively short in both directions, and response time is generally less than three seconds. Formerly known as the Airline Control Program/Transaction Processing Facility (ACP/TPF).

Transmission Control Protocol (TCP)

A communications 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 between hosts in packet-switched communications networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

Transmission Control Protocol/Internet Protocol (TCP/IP)

- 1) A combination of data-transmission protocols that provide end-to-end connections between applications over interconnected networks of different types.

2) A suite of transport and application protocols that run over the Internet Protocol. (GC) See also *Internet Protocol* and *Transmission Control Protocol*.

transparency

See *software transparency*.

TSO See *time sharing option*.

turbo processor

A faster multiprocessor that has six processors with common access to the main storage.

U

UFS UNIX filing system.

Ultra-SCSI

An enhanced small computer system interface.

unconfigure

To delete the configuration.

unit address

For zSeries, the address associated with a device on a given control unit. On ESCON interfaces, the unit address is the same as the device address. On OEMI interfaces, the unit address specifies a control unit and device pair on the interface.

unprotected volume

An AS/400 term that indicates that the AS/400 host recognizes the volume as an unprotected device, even though the storage resides on a RAID-formatted array and is, therefore, fault tolerant by definition. The data in an unprotected volume can be mirrored. Also referred to as an *unprotected device*.

upper-layer protocol

The layer of the Internet Protocol (IP) that supports one or more logical protocols (for example, a SCSI-command protocol and an ESA/390 command protocol). Refer to ANSI X3.230-199x.

UTC See *Coordinated Universal Time*.

V

virtual machine facility

A virtual data processing machine that appears to the user to be for the exclusive use of that user, but whose functions are accomplished by sharing the resources of

a shared data processing system. An alternate name for the VM/370 IBM operating system.

vital product data (VPD)

Information that uniquely defines the system, hardware, software, and microcode elements of a processing system.

VM

The root name of several IBM operating systems, such as VM/XA, VM/ESA, VM/CMS, and z/VM. See also *virtual machine facility*.

volume

For zSeries, the information recorded on a single unit of recording medium. Indirectly, it can refer to the unit of recording medium itself. On a nonremovable-medium storage device, the term can also indirectly refer to the storage device associated with the volume. When multiple volumes are stored on a single storage medium transparently to the program, the volumes can be referred to as logical volumes.

volume group

A collection of either physical or logical volumes.

volume label

A unique identifier that a user assigns to a logical volume.

VPD See *vital product data*.

VSE/ESA

An IBM operating system, the letters of which represent virtual storage extended/enterprise systems architecture.

W

weight distribution area

The area that is required to distribute the weight of the storage unit.

worldwide node name (WWNN)

A unique 64-bit identifier for a host that contains a fibre-channel port. See also *worldwide port name*.

worldwide port name (WWPN)

A unique 64-bit identifier associated with a fibre-channel adapter port. It is assigned in an implementation- and protocol-independent manner. See also *worldwide node name*.

write hit

A write operation in which the requested data is in the cache.

write penalty

The performance impact of a classical RAID-5 write operation.

WWNN

See *worldwide node name*.

WWPN

See *worldwide port name*.

X**xSeries**

The product name of an IBM e(logo)server product that emphasizes industry-standard server scalability and self-managing server technologies. It is the successor to the Netfinity family of servers.

Z**z/Architecture**

An IBM architecture for mainframe computers and peripherals. The IBM e(logo)server zSeries family of servers uses the z/Architecture architecture. It is the successor to the S/390 and 9672 family of servers. See also *iSeries*.

z/OS An operating system for the IBM e(logo)server product line that supports 64-bit real storage.

z/OS Global Mirror

A function of a storage server that assists a control program to maintain a consistent copy of a logical volume on another storage unit. All modifications of the primary logical volume by any attached host are presented in order to a single host. The host then makes these modifications on the secondary logical volume. This function was formerly called extended remote copy or XRC.

zSeries

An IBM e(logo)server family of servers that emphasizes near-zero downtime.

IBM enterprise servers based on z/Architecture.

zSeries storage

Storage arrays and logical volumes that are defined in the DS8000 as connected to zSeries servers.

Index

Numerics

524-byte sector, disk formatting 128

A

- accessibility features 159
 - keyboard 159
 - shortcut keys 159
- activating licenses 84
- adapters
 - attaching to open-systems host 12
- adding
 - users 91
- API, DS Open 9
- arrays
 - across loops 15
 - across loops overview 15
- arrays across loops 15
- availability features 10

B

backup and restore commands 133

C

- command-line interface 9
- commands
 - backup and restore 133
 - cpio 132
 - cplv 129
 - dump and restore 133
 - find 132
 - lsiv -l 131
 - lsiv -p rootvg 131
 - migratepv 131
 - migratepv -l 131
 - mkivcopy 130
 - scopy 137
 - xcopy 137
- comparing
 - Copy Services
 - features 20
- consistency groups
 - overview 122
- control unit (CU) 100
- Copy Services
 - across 2105 and 1750 107
 - across 2105 and 2107 107
 - disaster recovery 19, 106, 141
 - feature comparison 20
 - functions of 16, 105
 - Metro/Global Mirror 125
 - monitoring 16
 - overview 16, 105
 - remote
 - z/OS Global Mirror 124

- Copy Services 2105 domain
 - configuring 107
- copying
 - logical volume 129
- cpio commands 132
- cplv command 129
- creating
 - Metro Mirror
 - using a 2105 108
- CU (control unit) 100

D

- data migration
 - associated with physical volumes 131
 - backup and restore 137
 - backup and restore methods 133
 - between storage servers 138
 - copying a logical volume 129
 - direct copy 132
 - from eServer i5 400 134
 - from iSeries 134
 - from S/390 host systems 134
 - from Unix systems 135
 - from Windows host systems 137
 - from zSeries host systems 134
 - logical volume management software 129
 - migratepv commands 131
 - mirroring 130
 - moving workloads when migrating data between storage units 138
 - offloading workloads when migrating data between storage units 138
 - preparing for 127
 - reformatting existing disk drive modules 128
 - replacing existing storage
 - considerations 128
 - selecting method 23, 127, 128
- DDM
 - formatting 128
- defining
 - the logical control unit 103
- deleting
 - users 92
- description 112
- disaster recovery 106
 - failback 141
 - failover 141
 - fast restore 119
 - Global Mirror 146
 - Global Mirror planned scenario 148
 - Global Mirror scenario for unplanned 149
 - two-site solution 146
 - using Copy Services 19, 141
 - using the fast reverse restore feature 152
- disk formatting, 524-byte sector 128
- DS command-line interface 9
- DS Open API 9

- DS storage management
 - Linux graphical install 53
 - Linux graphical upgrade 72
 - Windows uninstall 42
- DS Storage Manager 9
- DS8000 series 1
 - Model 921 1, 2
 - Model 922 1, 2
 - Model 92E 1, 3
 - Model 9A2 1, 2
 - Model 9AE 1, 3
- DS8000 storage management
 - Linux silent mode install 64
 - Linux silent mode upgrade 80
 - Windows
 - silent mode install 37
 - silent mode upgrade 52
- DS8000 Storage Manager
 - firewall alert 25
 - hardware requirements 25
 - installation prerequisites 25
 - operating system support 25
 - Windows graphical install 26
 - Windows graphical upgrade 44
- DS8100 1
- DS8300 1
- DSFA, Disk Storage Feature Activation 84
- dump and restore commands 133

E

- error messages
 - generic SNMP 153
 - product-specific 154
 - storage unit 156
- ESCON
 - host systems 13
- Extended Remote Copy (XRC) (see z/OS Global Mirror) 124

F

- failback and failover 123
- failover
 - site switch 141, 143
- fast reverse 119
- fast reverse restore
 - advantages of using 152
- fast reverse, Global Mirror 152
- Feature Activation (DSFA), IBM Disk Storage 84
- fibre-channel
 - adapters 12
 - host systems 12
- FICON
 - host systems 14
- find commands 132
- FlashCopy
 - changing direction of 111
 - incremental 110
 - multiple relationships 113
 - overview 18, 110

- FlashCopy (*continued*)
 - persistent
 - overview 113
 - refresh target volume 114

G

- Global Copy
 - overview 117
- Global Mirror
 - description 118
 - disaster recovery for unplanned 149
 - disaster recovery without fast reverse 152
 - fast reverse restore 152
 - key terms defined 119
 - planned failover and failback scenario 148
 - setting up your environment 147
 - unplanned failover and failback 149
- guidelines
 - for defining the logical control unit 103

H

- host systems
 - fibre channel 12
 - S/390 and zSeries 13
 - supported by the storage unit 11

I

- IBM Disk Storage Feature Activation (DSFA) 84
- IBM TotalStorage Resiliency Family
 - Copy Services 106
 - introducing 106
- incremental FlashCopy
 - refreshing
 - FlashCopy target 114
- informational messages 156
- input/output configuration program (IOCP) 101
- IOCP
 - example 101
 - overview of the IOCP (input/output configuration program) 101

L

- LaunchPad, Windows install 26
- LCU
 - guidelines for defining the LCU (logical control unit) 103
- licensed functions 84
- licenses
 - Disk Storage Feature Activation (DSFA) 84
 - function authorization documents 84
- limitations
 - remote FlashCopy 114
- Linux
 - graphical install, DS storage management 53
 - graphical upgrade, DS storage management 72
 - silent mode install 64

Linux (*continued*)
 silent mode upgrade 80
logical control unit (LCU) 103
logical volume manager
 software 129
lslv -l commands 131
lsvg -p rootvg commands 131

M

management console
 multiple units 11, 87
 overview 11, 87
Metro Mirror
 overview
 write sequence 121
Metro/Global Mirror
 Copy Services 125
migratepv -l commands 131
migratepv commands 131
migrating data
 associated with physical volumes 131
 backup and restore 137
 backup and restore methods 133
 between storage servers 138
 copying a logical volume 129
 direct copy 132
 from eServer i5 400 134
 from iSeries 134
 from S/390 host systems 134
 from Unix systems 135
 from Windows host systems 137
 from zSeries host systems 134
 logical volume management software 129
 migratepv commands 131
 mirroring 130
 moving workloads when migrating data between
 storage units 138
 offloading workloads when migrating data between
 storage units 138
 preparing for 127
 reformatting existing disk drive modules 128
 replacing existing storage
 considerations 128
 selecting method 23, 127
MIH (missing-interrupt handler) setting 13
missing-interrupt handler (MIH) setting 13
mklvcopy command 130
modifying
 users 91
Multiple Device Manager (MDM), Copy Services 16
multiple relationship FlashCopy 113

O

obtaining activation codes 84
open-systems
 hosts
 fibre-channel 12
optional features codes
 copy services 99

optional features codes (*continued*)
 Parallel Access Volumes (PAVs) 99
overview
 consistency groups 122
 Copy Services
 monitor 16
 Multiple Device Manager (MDM) 16
 FlashCopy 18, 110
 Global Copy 117
 input/output configuration program (IOCP) 101
 Remote Mirror and Copy 116
 sessions management 121

P

parallel access volume (PAV)
 understanding static and dynamic 100
parallel access volumes for S/390 hosts 99
passwords
 use criteria 91
path
 connections 116
PAV (parallel access volumes) 99, 100
PAVs
 defining an LCU 103
 increasing the formatted 3390s 104
 IOCP program 101
 reconfiguring an LCU 104
 unit requirements 100
performing
 site switch
 failover 141, 143
persistent FlashCopy
 overview 113
physical configuration of DS8000
 battery assemblies 2, 3
 device adapters 2, 3
 disk enclosures 2, 3
 host adapters 2, 3
 I/O enclosures 2, 3
 Management Consoles 2
 power line disturbance (PLD) 3
problem states 154
processors 2
product-specific error messages 154

R

RAID
 RAID 10 overview 15
 RAID 5 overview 15
 RAID overview 15
reconfiguring an LCU 104
recovering 146
 with and without the fast reverse restore
 feature 152
refresh volume, FlashCopy 110
remote FlashCopy
 limitations 114
Remote Mirror and Copy 116
remote mirror for zSeries (see z/OS Global Mirror) 124

requirements, parallel access volumes 99
reversing
 FlashCopy relationship 111

S

S/390
 host systems 14
S/390 and zSeries
 host systems 13
 operating system 13
scenarios
 configuring
 Copy Services 2105 domain 107
scopy commands 137
service information messages
 S/390 systems 153
 zSeries systems 153
sessions, understanding 121
SNMP trap
 messages 153
software
 S/390 and zSeries 13
storage system logical partition
 overview 95
storage unit
 error messages 156
 host systems supported by 11
synchronous data mirroring 121
systems
 S/390 and zSeries host 13

T

Trademarks 163
troubleshooting
 generic SNMP trap messages 153
 informational messages 156
 problem states 154
 product-specific error messages 154
 service information messages 153
 storage unit error messages 156

U

uninstall DS storage management
 Windows 42
updating the I/O configuration 103
users
 adding 91
 deleting 92
 modifying 91
using
 fast reverse restore 152
using Copy Services 106
using FlashCopy
 with remote mirror and copy 112

W

Windows
 DS Storage Manager upgrade, graphical 44
 graphical install, DS8000 Storage Manager 26
 install, LaunchPad 26
 uninstall, DS storage management 42
Windows silent mode install 37
Windows silent mode upgrade 52

X

xcopy commands 137
XRC (Extended Remote Copy) (see z/OS Global Mirror) 124

Z

z/OS Global Mirror 124
zSeries
 host systems 14
 See S/390 and zSeries 13

Readers' comments — we would like to hear from you

IBM TotalStorage DS8000
User's Guide

Publication No. SC26-7623-03

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Cut or Fold
Along Line

Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department 61C
9032 South Rita Road
Tucson, Arizona 85775-4401



Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold
Along Line



Printed in USA

SC26-7623-03

