

IBM System Storage DS8000



User's Guide

Version 1 Release 4

IBM System Storage DS8000



User's Guide

Version 1 Release 4

Note:

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

This edition replaces SC26-7915-01 and all previous versions of SC26-7623.

© Copyright International Business Machines Corporation 2004, 2008. 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
-----------------------------------	-----------

Notices and publication information	xiii
--	-------------

Safety notices	xiii
Environmental notices	xiii
Product recycling and disposal	xiii
Battery return program	xiv
Conventions used in this guide.	xv
Related information	xvi
DS8000 series library	xvi
Other IBM publications	xvii
Ordering IBM publications	xxii
Web sites.	xxii
How to send your comments.	xxiii

Summary of Changes for SC26-7915-02 IBM System Storage DS8000 User's Guide	xxv
---	------------

Chapter 1. Introduction to IBM System Storage DS8000 series **1**

Overview of the DS8000 series models.	1
DS8100 (Model 921 and Model 931).	4
DS8300 (Models 922, 932, 9A2, and 9B2)	5
Model comparison	6
DS8000 physical footprint	8
DS8000 series performance features	10
DS8000 Interfaces	11
IBM System Storage DS Storage Manager	11
The DS command-line interface.	12
DS open application programming interface	13
IBM TotalStorage Productivity Center for Disk.	13
IBM TotalStorage Productivity Center for Replication	14
DS8000 hardware specifics	15
Storage unit structure	15
Host attachment overview	16
Processor memory	19
Overview of DS8000 power controls	19
Powering a storage unit on and off (real-time only)	20
Unit emergency power off (UEPO) switch	20
Rack operator panel	21
Data management features	21
RAID implementation	21
Arrays across loops.	23
Storage system logical partitions	23
IBM System Storage Resiliency Family	23
Comparison of licensed functions	29
Parallel access volumes and IBM HyperPAV	30

CKD extended address volumes - 3390 Model A	30
DS8000 series limitations	31
Planning data migration	32
Logical configuration overview.	32

Chapter 2. Installing the simulated DS8000 Storage Manager interface . . . **35**

Installing the DS Storage Manager on the Windows operating system using the graphical mode.	36
Installing the IBM System Storage DS Storage Manager on the Windows operating system using unattended (silent) mode	43
Removing the DS Storage Manager from a Windows operating system	50
Removing the DS Storage Manager from a Windows operating system using the unattended (silent) mode	51
Upgrading the DS Storage Manager on a Windows operating system using the graphical mode.	51
Upgrading the DS Storage Manager on the Windows operating system using unattended (silent) mode	56

Chapter 3. Installing the simulated IBM System Storage DS8000 Storage Manager on the Linux operating system. **57**

Installing the IBM System Storage DS Storage Manager on a Linux operating system using the graphical mode	57
Installing the IBM System Storage DS8000 Storage Manager on the Linux operating system in unattended (silent) mode	63
Removing the IBM System Storage DS Storage Manager from a Linux operating system.	70
Removing the DS8000 Storage Manager from a Linux operating system using the unattended (silent) mode	71
Upgrading the IBM System Storage DS Storage Manager on a Linux operating system	71
Upgrading the IBM System Storage DS8000 Storage Manager on the Linux operating system in unattended (silent) mode	76

Chapter 4. DS Storage Manager postinstallation instructions **79**

Internet browser support	79
Stop and start the DS Storage Manager Servers	80
Stopping and starting the DS Storage Manager servers on a Windows operating system.	80
Stopping and starting the DS Storage Manager servers on a Linux operating system	80
Starting the DS Storage Manager program	81

Starting the DS Storage Manager on a Windows operating system	81
Starting the DS8000 Storage Manager on a Linux operating system	82
Getting started with the DS Storage Manager	82
Activating licensed functions	83

Chapter 5. IBM System Storage Management Console 87

Logging onto the management console	87
Logging off the management console	88
Changing the management console password	88
Accessing the DS Storage Manager	88
Dual Hardware Management Consoles	89
Internet Protocol Conventions for IPv4 and IPv6	89

Chapter 6. Administering user accounts 91

User Groups	91
Adding user accounts	93
Modifying user accounts	94
Removing user accounts	94
Unlocking a user account	95
Unlocking an administrative password	95

Chapter 7. Storage system logical partition overview 97

DS8300 LPAR implementation	97
Variable logical partition implementation	98
Required hardware components to build a storage image	99
I/O resources	99
Processor and memory allocation	100
RIO-G interconnect separation	100

Chapter 8. Optional feature codes for the DS8000 101

Using parallel access volumes for System z and S/390 hosts	101
Parallel access volumes	102
PAV requirements for the DS8000	102
Overview of the input/output configuration program for the DS8000	103
Guidelines for defining the logical control unit for the DS8000	105
Reconfiguring an LCU with larger capacity 3390 volumes	106
Increasing the number of 3390s formatted within an LCU	106
IBM DS8000 Copy Services functions	107
FlashCopy	108
Remote mirror and copy	112
Metro/Global Mirror	118
Path connections	129
Copy Services for System z and S/390 hosts	130
Consistency groups	130
Failover and failback operations	132
Copy Services functions across a 2105 and 2107/242x	133

Copy Services functions with i5/OS	133
Adding a Copy Services 2105 domain to a storage complex	135
Creating a Metro Mirror volume pair between a 2107 /242x and a 2105	136
Creating remote FlashCopy transactions	138

Chapter 9. Migrating data to a storage unit. 141

How to select a data migration method	141
Considerations for replacing existing storage	142
Reformatting disk drive modules on open systems	142
Overview of the logical volume management software	143
Copying a complete logical volume	143
Using mirroring to migrate data	144
Using the migratepv command	145
Using the direct copy method	146
Using backup and restore methods to migrate data	147
Migrating data from an IBM System i host system to the DS8000	147
Migrating data from an IBM System z host system to the DS8000 storage unit	148
Data migration from a UNIX-based host system to a DS8000 storage unit	149
Migrating data to a DS8000 storage unit using volume management methods	150
Migrating data using the direct copy method	151
Migrating data using backup and restore methods	151
Migrating data from a Windows host system	151
Migrating data between storage units	151
Moving workloads when migrating data between storage units	152
Offloading workloads when migrating data between storage units	152

Chapter 10. Disaster recovery using Copy Services 153

Performing failover and failback operations (without Global Mirror)	154
Moving production to Site B after planned outages (failover)	154
Returning production to local site after planned outages	156
Moving production to Site B after unplanned outages (failover)	158

Appendix. Troubleshooting and analyzing problems. 159

Accessibility. 175

Notices 177

Terms and conditions	178
Trademarks	179
Electronic emission notices	180
Federal Communications Commission (FCC) statement	180

Industry Canada compliance statement	181
European Union EMC Directive conformance statement	181
Japanese Voluntary Control Council for Interference (VCCI) class A statement	182
Korean Ministry of Information and Communication (MIC) statement	183
Taiwan class A compliance statement	183

Taiwan Contact Information	183
--------------------------------------	-----

Glossary	185
.	185

Index	215
------------------------	------------

Figures

1. A base model (with front cover off) and its main components	2	8. DS8300 default LPAR resource allocation	98
2. An expansion model (with front cover off) and its main components	3	9. Storage image resource allocation in the processor complexes of the DS8300.	99
3. Maximum configuration for Model 921 or Model 931	4	10. Refreshing target volume — reverse restore	110
4. Configuration for Model 922, 932, 9A2, or 9B2 with two expansion models	6	11. Metro/Global Mirror configuration	119
5. DS8000 physical footprint. Dimensions are in centimeters (inches).	9	12. Metro/Global Mirror volume configuration	120
6. Rack operator panel.	21	13. Metro/Global Mirror configuration with master and subordinate storage units	121
7. Logical configuration sequence	34	14. Forming consistency groups according to specified intervals	124
		15. Incremental Resynchronization implementation	126

Tables

1.	DS8000 model comparison chart	6	4.	Failover and failback implementation	155
2.	Comparison of licensed functions	29	5.	Remote mirror and copy return codes	160
3.	User Group capabilities	92	6.	Copy Services suspension reason codes	163

About this guide

IBM System Storage™ 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 DS8000 Storage Manager on a Windows or Linux operating system
- Using the management console to access the DS8000 Storage Manager
- Using the optional features supported on the storage unit
- Migrating data to or between storage units
- Handling disaster recovery
- Analyzing problems

Note: For additional information on analyzing problems and troubleshooting, see “Troubleshooting” on page 159

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

Notices and publication information

This section contains information about safety notices that are used in this guide, environmental notices for this product, publication information, and information about sending your comments to IBM.

Safety notices

Complete this task 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 System Storage Solutions Safety Notices for IBM Versatile Storage Server and IBM System Storage Enterprise Storage Server, GC26-7229*.


Environmental notices

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

Product recycling and disposal

This unit contains recyclable materials.

This unit must be recycled or discarded according to applicable local and national regulations. IBM® encourages owners of information technology (IT) equipment to responsibly recycle their equipment when it is no longer needed. IBM offers a variety of product return programs and services in several countries to assist equipment owners in recycling their IT products. Information on IBM product recycling offerings can be found on IBM's Internet site at <http://www.ibm.com/ibm/environment/products/prp.shtml>.



Notice: This mark applies only to countries within the European Union (EU) and Norway.

Appliances are labeled in accordance with European Directive 2002/96/EC concerning waste electrical and electronic equipment (WEEE). The Directive determines the framework for the return and recycling of used appliances as applicable throughout the European Union. This label is applied to various products to indicate that the product is not to be thrown away, but rather reclaimed upon end of life per this Directive.


In accordance with the European WEEE Directive, electrical and electronic equipment (EEE) is to be collected separately and to be reused, recycled, or recovered at end of life. Users of EEE with the WEEE marking per Annex IV of the WEEE Directive, as shown above, must not dispose of end of life EEE as unsorted municipal waste, but use the collection framework available to customers for the return, recycling and recovery of WEEE. Customer participation is important to minimize any potential effects of EEE on the environment and human health due to the potential presence of hazardous substances in EEE. For proper collection and treatment, contact your local IBM representative.

Battery return program

This product may contain a sealed lead acid, nickel cadmium, nickel metal hydride, lithium, or lithium ion battery. Consult your user manual or service manual for specific battery information. The battery must be recycled or disposed of properly. Recycling facilities may not be available in your area. For information on disposal of batteries outside the United States, go to <http://www.ibm.com/ibm/environment/products/index.shtml> or contact your local waste disposal facility.

In the United States, IBM has established a return process for reuse, recycling, or proper disposal of used IBM sealed lead acid, nickel cadmium, nickel metal hydride, and other battery packs from IBM Equipment. For information on proper disposal of these batteries, contact IBM at 1-800-426-4333. Please have the IBM part number listed on the battery available prior to your call.

In the Netherlands the following applies:



For Taiwan:



Please recycle batteries.

廢電池請回收

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

This section provides information about locating and ordering DS8000™ publications.

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

- The publications that make up the IBM® System Storage™ 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 xxii for information about how to order publications in the IBM System Storage DS8000 series publication library. See “How to send your comments” on page xxiii for information about how to send comments about the publications.

DS8000 series library

These customer publications make up the DS8000 series library.

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-8825. These publications are also available as PDF files on the following Support Web site:

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

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

Title	Description	Order Number
<i>IBM System Storage 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.	SC26-7916
<i>IBM System Storage DS8000: Host Systems Attachment Guide</i>	This guide provides information about attaching hosts to the DS8000 storage unit. The DS8000 provides a variety of host attachments so that you can consolidate storage capacity and workloads for open-systems hosts and System z™ or S/390® hosts.	SC26-7917
<i>IBM System Storage 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.	GC35-0515
<i>IBM System Storage 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 single devices that consolidate redundant paths for 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.	GC27-2122
<i>IBM System Storage DS8000: User's Guide</i>	This guide provides instructions for setting up and operating the DS8000 and for analyzing problems.	SC26-7915
<i>IBM System Storage DS Application Programming Interface Reference</i>	This publication provides reference information for the IBM System Storage DS application programming interface (API) and provides instructions for installing the Common Information Model Agent, which implements the API.	GC35-0516

Title	Description	Order Number
<i>IBM System Storage DS8000 Messages Reference</i>	This publication provides explanations of error, information, and warning messages that are issued from the DS8000 user interfaces.	GC26-7914

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 xxii 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 TotalStorage DS8000 Series: Implementation</i>	This publication, from the IBM International Technical Support Organization, helps you plan, install, and configure the IBM System Storage DS8000 series. It can help you plan, design, and implement a new installation or migrate from an existing one. This publication also discusses the architecture and components of the DS8000 and provides useful information for anyone who will manage the DS8000.	SG24-6786
<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 host system. Copy Services functions include peer-to-peer remote copy (PPRC), extended remote copy (XRC), 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>IBM System Storage Copy Services and System i5: A Guide to Planning and Implementation</i>	This publication, from the IBM International Technical Support Organization, tells you how to plan for and implement Copy Services on the IBM System i5™ platform using the IBM System Storage DS models. This publication includes examples for creating an IBM FlashCopy environment that you can use for offline backup or testing. It also includes examples for setting up the two availability products of Copy Services for disaster recovery: Metro Mirror and Global Mirror. For additional information on Copy Services, see the IBM International Technical Support Organization publications <i>iSeries and IBM TotalStorage: A Guide to Implementing External Disk on eServer i5</i> (SG24-7120) and <i>IBM System Storage DS8000 Series: Architecture and Implementation</i> (SG24-6786).	SG24-7103

Title	Description	Order Number
<i>IBM System Storage DS8000 Series: Copy Services in Open Environments</i>	This publication, from the IBM International Technical Support Organization, describes the Copy Services feature, which can provide immediate copies of data and make the data available to other applications or systems. The Copy Series feature provides local or remote copies for data retention and disaster recovery. Copy Services functions that are described include Metro Mirror, Global Mirror and Global Copy functions. This publication also describes the 3-site Metro/Global Mirror with Incremental Resync feature and introduces the TotalStorage Productivity Center for Replication.	SG24-6788
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
Open-systems hosts		
<i>IBM System Storage DS8000 Series: Copy Services in Open Environments</i>	This publication, from the IBM International Technical Support Organization, describes the Copy Services feature, which can provide immediate copies of data and make the data available to other applications or systems. The Copy Series feature provides local or remote copies for data retention and disaster recovery. Copy Services functions that are described include Metro Mirror, Global Mirror and Global Copy functions. This publication also describes the 3-site Metro/Global Mirror with Incremental Resync feature and introduces the TotalStorage Productivity Center for Replication.	SG24-6788
System z and S/390 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

Title	Description	Order Number
<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 System z 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>IBM System Storage DS8000 series: Copy Services with IBM System z</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on a DS8000 in a System z environment. This publication provides the necessary details to implement and control Copy Services functions. It also includes tips for maximizing effective installation and operation. Because IBM Advanced Copy Services functions can be used across DS8000 series, DS6000 series, and some models of the ESS family, this publication also discusses interoperability considerations.	SG24-6787
<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 the ResourceLink Web site at http://www.ibm.com/servers/resourcelink/ .	SB10-7029
<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 the ResourceLink Web site at http://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

Title	Description	Order Number
<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		
<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
Seascope® family		

Title	Description	Order Number
<i>IBM TotalStorage DS8000 Series: Implementation</i>	This publication, from the IBM International Technical Support Organization, helps you plan, install, and configure the IBM System Storage DS8000 series. It can help you plan, design, and implement a new installation or migrate from an existing one. This publication also discusses the architecture and components of the DS8000 and provides useful information for anyone who will manage the DS8000.	SG24-6786
<i>IBM TotalStorage DS8000 Series: 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 DS8000 series to ensure optimum performance.	SG24-7146
<i>Migrating to IBM System Storage DS8000</i>	This publication, from the IBM International Technical Support Organization, provides data migration information when the DS8000 is the target system. This publication also reviews products and describes IBM data migration services. After it describes appliance-based, storage-based, and host-based techniques, it provides migration scenarios that use various products and operating platforms.	SG24-7432
<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		
<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>Migrating to IBM System Storage DS8000</i>	This publication, from the IBM International Technical Support Organization, provides data migration information when the DS8000 is the target system. This publication also reviews products and describes IBM data migration services. After it describes appliance-based, storage-based, and host-based techniques, it provides migration scenarios that use various products and operating platforms.	SG24-7432
<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
System Storage Productivity Center		
<i>IBM System Storage Productivity Center Introduction and Planning Guide</i>	This publication introduces the IBM System Storage Productivity Center hardware and software.	SC23-8824
<i>IBM System Storage Productivity Center Hardware Installation and Configuration Guide</i>	This publication describes how to install and configure the IBM System Storage Productivity Center hardware.	SC23-8822

Title	Description	Order Number
<i>IBM System Storage Productivity Center Software Installation and User's Guide</i>	This publication describes how to install and use the IBM System Storage Productivity Center software.	SC23-8823

Ordering IBM publications

You can order copies of IBM publications using the IBM publications center.

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.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi>

Note: Open the Web site in a new browser window by right clicking on the link and selecting "Open in New Window."

Web sites

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

Type of Storage Information	Web Site
Concurrent Copy for System z and S/390 host systems	http://www.storage.ibm.com/software/sms/sdm/
DS command-line interface (DS CLI)	http://www-1.ibm.com/support/docview.wss?rs=1113&context=HW2B2&dc=DB500&uid=ssg1S1002949&loc=en_US&cs=utf-8&lang=en See Section 3 .
DS8000 Information Center	http://publib.boulder.ibm.com/infocenter/ds8000ic/index.jsp
DS8000 series publications	http://publib.boulder.ibm.com/infocenter/ds8000ic/index.jsp Click Publications .
FlashCopy for System z and S/390 host systems	http://www.storage.ibm.com/software/sms/sdm/
Host system models, operating systems, adapters, and switches that the DS8000 series supports	http://www.ibm.com/servers/storage/disk/ds8000/ Click Interoperability matrix . http://www.ibm.com/systems/support/storage/config/ssic/ Click New search .
IBM Disk Storage Feature Activation (DSFA)	http://www.ibm.com/storage/dsfa

Type of Storage Information	Web Site
IBM storage products	http://www.storage.ibm.com/
IBM System Storage 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 System z and S/390 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/
Subsystem Device Driver (SDD)	http://www.ibm.com/systems/support/storage/software/sdd
Technical notes and product tips	http://www.ibm.com/servers/storage/support/disk/ds8100/ Click Technical notes on the Troubleshooting tab.
z/OS Global Mirror (formerly XRC) for System z and S/390 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 SC26-7915-02 IBM System Storage 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. **Note: In this edition, the brand “IBM TotalStorage” has been replaced with “IBM System Storage”.**

New Information

As of Release 3 and all follow on releases, the stand-alone version of the DS Storage Manager can only be installed as the simulated DS Storage Manager and will only operate in Simulated-mode. Real-time mode is no longer supported or included for the stand-alone version.

If your installation is at a release **prior** to Release 3, you can still access the real-time DS Storage Manager with your remote browser by pointing your browser at the hardware management console (HMC).

If your installation is at Release 3, the real-time DS Storage Manager is **only** accessible through the console at the HMC or through the IBM System Storage Productivity Center. Accessing the real-time DS Storage Manager at the HMC or through the IBM System Storage Productivity Center helps to eliminate confusion when you use the simulated DS Storage Manager.

Here are some things to consider as a result of this change:

- You do not need to perform a stand-alone installation to use the real-time DS Storage Manager.
- You must manually enter the URL for the DS Storage Manager into the browser on the HMC and create your own bookmarks, or you must provide the IBM System Storage Productivity Center Element Manager with the information for the HMC to access the real-time DS Storage Manager.
- You do not need to install the stand-alone simulated DS Storage Manager unless you intend to use it.
- You can import data from the HMC to a stand-alone simulated DS Storage manager using the Import action on the simulated DS Storage Manager.
- You can use the simulated DS Storage Manager to create an initial configuration file that can be applied to the HMC.

Attention: The simulated DS Storage Manager is designed to assist you in creating an initial configuration for your system. Do not attempt to upload a configuration file if your system has already been configured.

The following new information is introduced in chapter 1:

- New DS8000 series models that are available. The new base models are Models 931, 932, and 9B2. See “Overview of the DS8000 series models” on page 1 for more information.
- The IBM TotalStorage Productivity Center for Disk. See “IBM TotalStorage Productivity Center for Disk” on page 13 for more information.

- The IBM TotalStorage Productivity Center for Replication Manager. See “IBM TotalStorage Productivity Center for Replication” on page 14 for more information.
- The Metro/Global Mirror function for Copy Services. See “Metro/Global Mirror” on page 118
- Chapter One was updated to include the following topics:
 - Overview of DS8000 power controls. See “Overview of DS8000 power controls” on page 19 for more information.
 - Powering a storage unit on and off. See “Powering a storage unit on and off (real-time only)” on page 20 for more information.
 - Unit emergency power off (UEPO) switch. See “Unit emergency power off (UEPO) switch” on page 20 for more information.
 - Rack operator panel. See “Rack operator panel” on page 21 for more information.
- Browser support for Mozilla Firefox 1.5 or 2.0. See Chapter 2, “Installing the simulated DS8000 Storage Manager interface,” on page 35 for more information.
- Specific trap 221 - Space Efficient Repository or Over-provisioned Volume has reached a warning watermark, was added to the troubleshooting section. See “Troubleshooting” on page 159 for more information.
- Support for Internet Protocol Version 6 (IPv6) was added. See “Internet Protocol Conventions for IPv4 and IPv6” on page 89 for more information.
- Support for RAID 6 was added. See “RAID 6 overview” on page 22 for more information.
- Support for larger volumes on z/Series was added through the use of CKD Extended Address Volumes - 3390 Model A. See “CKD extended address volumes - 3390 Model A” on page 30
- Support for a variable logical partition (LPAR) was added for DS8300 model 2107-9B2. See “Variable logical partition implementation” on page 98 for more information.
- Support for Internet Protocol Version 6 (IPv6) TPC on IBM TotalStorage Productivity Center for Replication Manager was added. See “IBM TotalStorage Productivity Center for Replication” on page 14 for more information.
-

Changed Information

The following information was updated or changed:

- Chapters 2, 3, and 4 have updated instructions for installing and upgrading the DS simulated Storage Manager.
- Chapter 6 has updated information for user group definitions and unlocking an administrative password.
- Chapters 8 and 10 were reorganized.
- Information regarding Global Mirror disaster recovery processing was removed from this publication and is now published in the *IBM System Storage DS8000 Command-Line Interface User's Guide*.

Chapter 1. Introduction to IBM System Storage DS8000 series

IBM System Storage DS8000 series is a high-performance, high-capacity series of disk storage that is designed to support continuous operations. DS8000 series models (machine type 2107/242x) 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. You can attach a maximum of four DS8000 series models to each management console. 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 information in this section to familiarize yourself with the DS8000 series.

You can also view the DS8000 series e-learning overview from the DS8000 Storage Manager Welcome page or at the following Web site:

<http://www-1.ibm.com/support/docview.wss?rs=1113&uid=ssg1S7001199>

The e-learning overview provides an animated presentation about installation and configuration, service and support, and management tasks.

Overview of the DS8000 series 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 (Models 921 and 931) features a dual two-way processor complex and support for one expansion model.

DS8300

The DS8300 (Models 922, 932, 9A2, and 9B2) features a dual four-way processor complex and support for up to four expansion models. Models 9A2 and 9B2 support two IBM System Storage System logical partitions (LPARs) in one storage unit.

DS8000 expansion models (Models 92E, 9AE) can be attached to a base model to expand its capabilities:

921 and 931 (base) can attach to 92E (expansion)

922 and 932 (base) can attach to 92E (expansion)

9A2 and 9B2 (base) can attach to 9AE (expansion)

Note: Models 921, 922, and 9A2 apply only to hardware machine type 2107. These models have been withdrawn from marketing. Although you can no longer order these hardware models directly from IBM, you can order add-on features or licensed functions for your existing 2107 hardware models. You can obtain the hardware models on an as-available basis through IBM Business Partners.

Figure 1 provides a high-level view of the components of a base model (Models 921, 922, 931, 932, 9A2, 9B2).




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. Models 921 and 931 contain two batteries. Models 922, 932, 9A2, and 9B2 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*, 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 Model 921 and 931 processor enclosures have 2-way processors. Processor enclosures on Models 922, 932, 9A2, and 9B2 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 provides a high-level view of the components of an expansion model (Models 92E, 9AE).




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 model contains I/O enclosures or the extended power line disturbance (PLD) feature, two or 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 *adapters* if they are the first expansion models that are attached to a Model 922, 932, 9A2, or 9B2. The second, third, or fourth expansion model in a 922, 932, 9A2, or 9B2 configuration cannot have I/O enclosures and adapters, nor can any expansion model that is attached to a Model 921 or 931. If the expansion model contains I/O enclosures, the enclosures provide connectivity between the adapters and the processors. The adapters contained in the I/O enclosures can be either device or host adapters.

DS8100 (Model 921 and Model 931)

The IBM System Storage DS8100, which is Model 921 and Model 931, offers many features.

The DS8100 provides the following features:

- Dual 2-way processor complex
- Up to 128 disk drives, for a maximum capacity of 64 TB
- Up to 128 GB of processor memory (cache)
- Up to 16 fibre channel/FICON or ESCON host adapters

The DS8100 base models can support one expansion model. With one expansion model, you can expand the capacity of the Model 921 or Model 931 with up to 384 disk drives, for a maximum capacity of 192 TB.

Note: IBM service representatives can upgrade a Model 921 or Model 931 in the field when you order a model conversion to a Model 922, Model 932, Model 9A2, or Model 9B2. There are no available model conversions from Models 92x/9Ax to Models 93x/9Bx.

Figure 3 shows the maximum configuration of a Model 921 or Model 931, which is the 921 or the 931 base model plus one 92E expansion model.

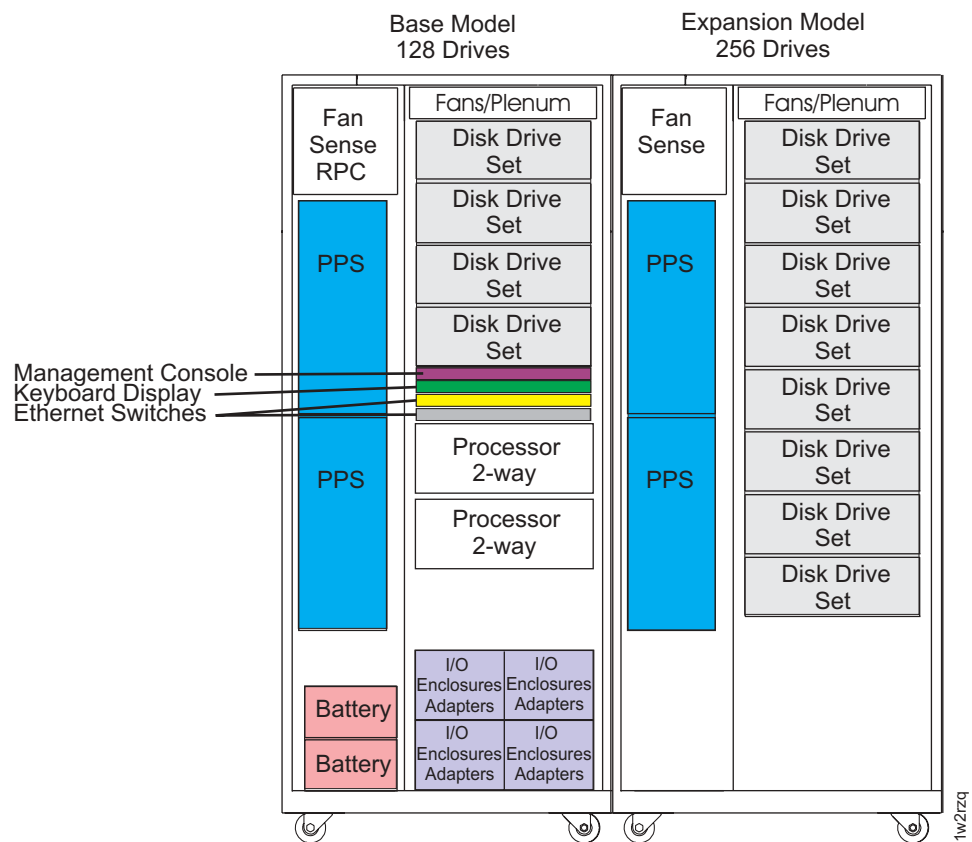


Figure 3. Maximum configuration for Model 921 or Model 931

DS8300 (Models 922, 932, 9A2, and 9B2)

IBM System Storage DS8300 models (Models 922, 932, 9A2, and 9B2) offer higher performance and capacity than the DS8100. Models 9A2 and 9B2 also enable you to create two storage system LPARs (images) within the same storage unit.

All DS8300 models offer the following features:

- Dual 4-way processor complex
- Up to 128 disk drives, for a maximum capacity of 64 TB
- Up to 256 GB of processor memory (cache)
- Up to 16 fibre-channel/FICON or ESCON host adapters

The DS8300 base models can support up to four expansion models. With expansion models, you can expand the Models 922, 932, 9A2, and 9B2 as follows:

- With one expansion model, you can support the following expanded capacity and number of adapters:
 - Up to 384 disk drives, for a maximum capacity of 192 TB
 - Up to 32 fibre-channel/FICON or ESCON host adapters
- With two expansion models, you can support the following expanded capacity:
 - Up to 640 disk drives, for a maximum capacity of 320 TB
- With three expansion models, you can support the following expanded capacity:
 - Up to 896 disk drives, for a maximum capacity of 448 TB
- With four expansion models, you can support the following expanded capacity:
 - Up to 1024 disk drives, for a maximum capacity of 512 TB

Figure 4 on page 6 shows the configuration of a Model 922, 932, 9A2, or 9B2 with two expansion models. Each of these models can attach up to four expansion models. Models 922 and 932 can attach to the 92E expansion models. Models 9A2 and 9B2 can attach to the 9AE expansion models.

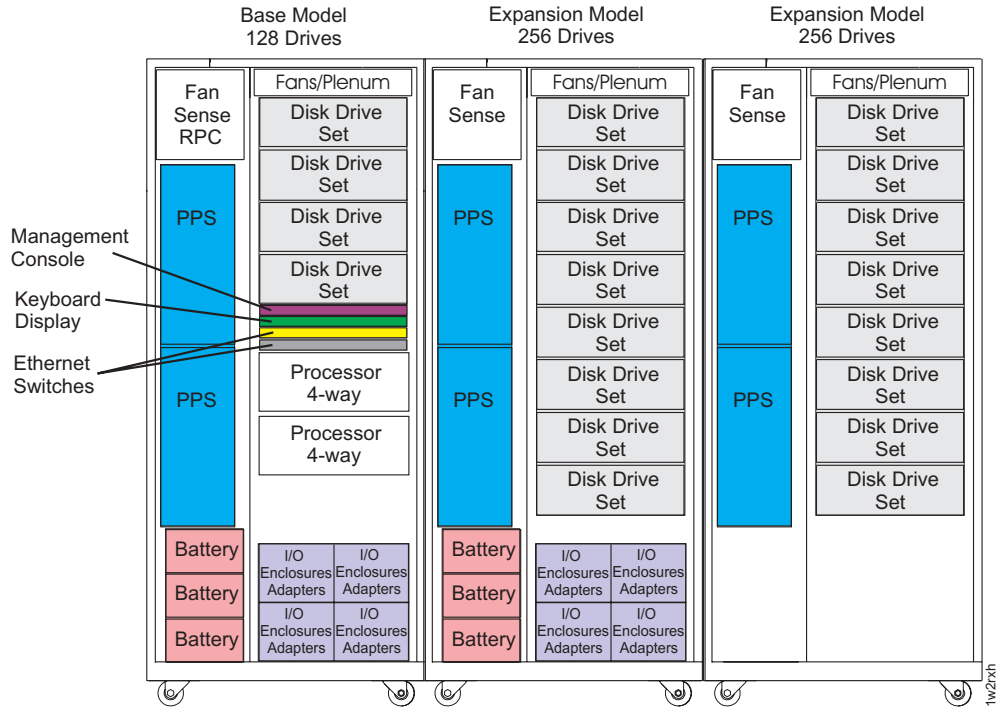


Figure 4. Configuration for Model 922, 932, 9A2, or 9B2 with two expansion models

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 921/931	1 image	None	2-way	Up to 128 disks (maximum capacity of 64 TB)	Up to 128 GB	Up to 16 host adapters
		With 1 expansion model (92E)	2-way	Up to 384 disks (maximum capacity of 192 TB)	Up to 128 GB	Up to 16 host adapters

Table 1. DS8000 model comparison chart (continued)

Base model	Images (logical partitions)	Expansion models	Processor type	Disk drives	Processor memory	Host adapters
Model 922/932	1 image	None	4-way	Up to 128 disks (maximum capacity of 64 TB)	Up to 256 GB	Up to 16 host adapters
	With 1 expansion model (92E)	With 1 expansion model (92E)	4-way	Up to 384 disks (maximum capacity of 192 TB)	Up to 256 GB	Up to 32 host adapters
		With 2 expansion models (92E)	4-way	Up to 640 disks (maximum capacity of 320 TB)	Up to 256 GB	Up to 32 host adapters
		With 3 expansion models (92E)	4-way	Up to 896 disks (maximum capacity of 448 TB)	Up to 256 GB	Up to 32 host adapters
		With 4 expansion models (92E)	4-way	Up to 1024 disks (maximum capacity of 512 TB)	Up to 256 GB	Up to 32 host adapters
Model 9A2/9B2	2 images	None	4-way	Up to 128 disks (maximum capacity of 64 TB)	Up to 256 GB	Up to 16 host adapters
	With 1 expansion model (9AE)	With 1 expansion model (9AE)	4-way	Up to 384 disks (maximum capacity of 192 TB)	Up to 256 GB	Up to 32 host adapters
		With 2 expansion models (9AE)	4-way	Up to 640 disks (maximum capacity of 320 TB)	Up to 256 GB	Up to 32 host adapters
		With 3 expansion models (9AE)	4-way	Up to 896 disks (maximum capacity of 448 TB)	Up to 256 GB	Up to 32 host adapters
		With 4 expansion models (9AE)	4-way	Up to 1024 disks (maximum capacity of 512 TB)	Up to 256 GB	Up to 32 host adapters

Note: IBM service representatives can perform the conversion in the field when you order one of the following model conversions:

- Model 921 to 922
- Model 921 to 9A2 (will be first converted to a 922, and then to a 9A2)
- Model 931 to 932
- Model 931 to 9B2 (will be first converted to a 932, and then to a 9B2)
- Model 9A2 to 922
- Model 922 to 9A2
- Model 932 to 9B2
- Model 9B2 to 932
- Model 92E to 9AE
- Model 9AE to 92E

DS8000 physical footprint

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

Figure 5 on page 9 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 series 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:

Storage pool striping (rotate extents)

Storage pool striping is now supported on the DS8000 series, providing improved performance. The storage pool striping function stripes new volumes across all ranks of an extent pool. The striped volume layout reduces administration that is required to balance system loads. With storage pool striping support, the system automatically performs close to highest efficiency, requiring little or no administration. The effectiveness of performance management tools is also enhanced, because imbalances tend to occur as isolated problems. When performance administration is required, it is applied more precisely. Storage pool striping can be managed and configured via the DS Storage Manager, DS CLI, and DS Open API. The storage pool striping function is provided with the DS8000 series at no additional charge.

POWER5+ processor technology

The DS8000 series features IBM POWER5+ server technology. Compared with the current IBM POWER5 processor, the POWER5+ processor allows up to a 15% performance improvement in I/O operations per second in transaction processing workload environments. The POWER5+ processor is included only in the 93x/9Bx models, and it is an optional feature in the 92x/9Ax models.

Industry standard fibre-channel disk drives

The DS8000 series offers a selection of fibre-channel disk drives, including 300 GB fibre channel drives and 500 GB fibre-channel ATA (FATA) disk drives that allow a DS8100 to scale up to a capacity of 192 TB and a DS8300 to scale up to a capacity of 512 TB. The 500 GB FATA disk drives are lower cost disk drives that are designed for applications that require disk drive duty cycles that are equal to or less than 20%. If you use the FATA disk drives for applications that require duty cycles greater than 20%, your application is likely to experience limited array throughput and reduced performance. The 500 GB FATA disk drives are designed to support applications that use fixed-content or reference-only data, applications that are used for data archival operations, or near-line applications that require large amounts of data. Intermix of the FATA disk drive sets and fibre-channel disk drive sets is not supported within a disk enclosure or at creation. However, intermix on the same DA loop of enclosures with different drive types is supported if the drives within each enclosure are the same.

Four-port fibre-channel/FICON adapters

The DS8000 series offers enhanced connectivity with four-port fibre-channel/FICON host adapters. Four GB fibre-channel/FICON host support offers improvement that is up to 50 percent of the MB/second throughput performance for a single port. This improved performance enables cost savings by reducing the number of required host ports. These 4 GB host adapters are offered in longwave and shortwave. They auto-negotiate to either 4 GB or 2 GB link speeds. This flexibility provides

benefits that are offered by higher performance 4 GB SAN-based solutions, while maintaining compatibility with existing 2 GB infrastructures.

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 nondisruptive 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. 500 GB, 7 200 rpm fibre-channel ATA (FATA) Standby Capacity on Demand (Standby CoD) disk drive sets are available as part of the Standby CoD offering.

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 series 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.

For the latest version of the online Information Center, go to <http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp>

DS8000 Interfaces

This section describes the following interfaces:

- IBM System Storage DS Storage Manager
- IBM System Storage DS Command-Line Interface (CLI)
- IBM System Storage DS Open application programming interface
- IBM TotalStorage Productivity Center for Disk
- IBM TotalStorage Productivity for Replication Manager

IBM System Storage DS Storage Manager

The IBM System Storage DS Storage Manager is an interface that is used to perform logical configurations and Copy Services management functions.

The real-time DS Storage Manager software is installed on a computer (management console) that is provided with the DS8000 series. The real-time DS Storage Manager can be accessed using the management console or through the IBM System Storage Productivity Center. The real-time DS Storage Manager supports the following components:

Real-time configuration

This component is preinstalled on your management console. It provides you with 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.

Copy services

This component allows you to process copy services functions.

You can install the simulated DS Storage Manager using the GUI (graphical mode) or using the unattended (silent) mode for the supported operating systems. Use the unattended (silent) mode of installation to customize a response file and issue a command from a command prompt window. The response file is a template on the IBM System Storage Storage Manager CD. The simulated DS Storage Manager can be accessed from any location that has network access using a Web browser. It provides a simulated configuration option that is available with the applicable license management code.

You can install this component on your PC which provides the ability to create logical configurations. After creating the configuration, you can save it and then apply it to a network-attached storage unit at a later time. This simulated configuration tool is installed on your server and can be used for the configuration of a DS8000 series at initial installation.

Attention: The simulated DS Storage Manager is designed to assist you in creating an initial configuration for your system. Do not attempt to upload a configuration file if your system has already been configured.

The DS command-line interface

The IBM System Storage DS command-line interface (CLI) allows a user to create, delete, modify, and view Copy Services functions and the logical configuration of a storage unit. These tasks can be performed either interactively, in batch processes (operating system shell scripts), or in DS CLI script files. A DS CLI script file is a text file that contains one or more DS CLI commands and can be issued as a single command.

The DS CLI provides a full-function set of commands to manage logical configurations and Copy Services configurations on DS8000, DS6000™, and Enterprise Storage Server (ESS) storage units. The DS CLI can be installed on and is supported in many different environments (for example, AIX®, i5/OS®, HP-UX, Linux®, Netware, OpenVMS, Solaris, and Windows®). Batch processes and task schedulers (for example, cron), can be used to automate many DS CLI common tasks (for example, data replication).

Notes:

1. For DS8000, you can have a maximum of 256 clients connected to a single storage unit server at the same time. Clients include all DS Storage Managers, DS command-line interfaces, DS open application interfaces, and IBM TotalStorage Productivity Center for Replication sessions. However, you must not simultaneously start more than 100 client sessions including DS CLI sessions. Starting more than 100 sessions simultaneously can result in connection problems.
2. You cannot install the DS8000 series DS CLI on a Windows 64-bit operating system.

DS open application programming interface

The IBM System Storage 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 RAID 5, RAID 6, and RAID 10 volume spaces. The DS Open API supports these activities through the use of the Storage Management Initiative Specification (SMI-S), as defined by the Storage Networking Industry Association (SNIA).

The DS Open API helps integrate 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 System Storage DS Storage Manager web-based interface and the DS command-line interface.

You must implement the DS Open API through the IBM System Storage 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.

The DS Open API supports the IBM System Storage DS8000 and the IBM System Storage DS6000 series, and the IBM TotalStorage Enterprise Storage Server. It is available for the AIX, Linux, and Windows operating system environments and must be used on storage units that have fibre-channel ports.

IBM TotalStorage Productivity Center for Disk

IBM TotalStorage Productivity Center for Disk (Disk Manager) is a storage management product that is designed for customers who want to reduce the complexities and costs of storage management, including the management of SAN-based storage.

The IBM TotalStorage Productivity Center allows you to manage storage area networks and heterogeneous storage from a single console. It lets you manage network storage components that are based on the Storage Management Initiative Specification, including

- IBM SAN Volume Controller
- IBM TotalStorage Enterprise Storage Server (ESS)
- IBM System Storage disk storage units (DS4000, DS6000, and DS8000 series)

Disk Manager can do the following performance functions:

- Collect and store performance data and provides alerts
- Generate both standard and custom performance reports with the ability to graph data from those reports
- Help optimize storage allocation
- Provide volume contention analysis and recommend intelligent placement of new volumes

Through the use of data collection, threshold settings, and performance reports, performance can be monitored for the ESS, DS4000, DS6000, and DS8000, SAN

Volume Controller, and other storage subsystems that support the SMI-S block service performance subprofile. The performance functions start with the data collection task, which captures performance statistics for the devices and stores the data in the database.

You can set thresholds for certain performance metrics depending on the type of device. Threshold checking is performed during data collection, and when performance is outside the specified bounds, alerts can be generated.

After performance data has been collected, you can configure Disk Manager to present graphical or text reports on the historical performance behavior of specified devices, or of specific internal components of these devices. The performance reports provide information on the performance metrics and can display performance data in graphical form.

For more information, visit the Multiple Device Manager Information Center from the IBM Publications Web site using the following URL:

<http://publib.boulder.ibm.com/infocenter/tsmdmv11/index.jsp>

IBM TotalStorage Productivity Center for Replication

IBM TotalStorage Productivity Center for Replication provides Copy Services functions such as the remote mirror and copy functions (Metro Mirror and Global Mirror) and the point-in-time function (FlashCopy).

TotalStorage Productivity Center for Replication provides a graphical interface that you can use for configuring and managing Copy Services functions across the DS8000, DS6000, and Enterprise Storage Server (ESS) storage units. These data-copy services maintain consistent copies of data on source volumes that are managed by Replication Manager.

IBM TotalStorage Productivity Center for Replication V3.1 for FlashCopy, Metro Mirror, and Global Mirror support provided automation of administration and configuration of these services, operational control (starting, suspending, resuming), Copy Services tasks, and monitoring and managing of copy sessions.

TotalStorage Productivity Center for Replication is an option of the TotalStorage Productivity Center for Replication software program. If you are licensed for Copy Services functions and have the required Ethernet adapter pair (feature code 180x) that supports the communication between the storage units and Replication Manager, you can use TotalStorage Productivity Center for Replication to manage your data copy environment.

Notes:

1. TotalStorage Productivity Center for Replication operations can now be performed using the DS8000 hardware management console (HMC).
2. The use of Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6) are both supported through the HMC ports.
3. Ethernet ports on the DS8000 server only support the use of IPv4.

For more information, visit the Multiple Device Manager Information Center from the IBM Publications Web site using the following URL:

<http://publib.boulder.ibm.com/infocenter/tsmdmv11/index.jsp>

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, four 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 System Storage 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.

Host attachment overview

The DS8000 storage unit provides a variety of host attachments so that you can consolidate storage capacity and workloads for open-systems hosts and System z or S/390 hosts.

The DS8000 provides extensive connectivity using fibre-channel or ESCON adapters across a broad range of server environments. 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

The DS8100 Models 921 and 931 support a maximum of 16 host adapters and 4 device adapter pairs, and the DS8300 Models 922, 932, 9A2, and 9B2 support a maximum of 32 host adapters and 8 device adapter pairs.

To reduce the possibility of interactions between system adapters in switched configurations, you can establish zones for your fibre-channel attachments. You can minimize contention issues if the zones contain a single port that is attached to a system adapter with the desired number of ports that are attached to the storage unit.

To establish zones, use either of the following 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 host 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.

Open-systems host attachment 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, 2 Gbps, or 4 Gbps full-duplex, serial communications technology to interconnect I/O devices and host systems that are separated by tens of kilometers.

The IBM System Storage DS8000 series supports SAN connections of up to 2 Gbps with 2 Gb host adapters, and up to 4 Gbps with 4 Gb host adapters. The DS8000 series negotiates automatically and determines whether it is best to run at a 1 Gbps, 2 Gbps, or 4 Gbps link speed. The IBM System Storage DS8000 series detects and operates at the greatest available link speed that is shared by both sides of the system.

Note: 4 Gbps link speeds are only available with 4 Gbps fibre-channel adapter cards.

Fibre-channel technology 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 overview for the DS8000 series

Each storage unit fibre-channel adapter has four ports, and each port has a unique worldwide port name (WWPN). You can configure a port to operate with the SCSI-FCP upper-layer protocol using the DS Storage Manager or the DS CLI. You can add fibre-channel shortwave and longwave adapters to a DS8000 model.

For details on the host systems that support fibre-channel adapters, see the DS8000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/> or the System Storage Interoperation Center (SSIC) at <http://www.ibm.com/systems/support/storage/config/ssic/>.

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

- A maximum of 64 host adapter ports on DS8100 Models 921 and 931 and a maximum of 128 host ports on DS8300 Models 922, 932, 9A2, and 9B2
- A maximum of 506 logins per fibre-channel port which includes host ports and PPRC target and initiator ports
- Access to 63700 LUNs per target (one target per host adapter), depending on host type
- Either arbitrated loop, switched-fabric, or point-to-point topologies

ESCON-attached System z and S/390 hosts overview

This section describes how you can attach the DS8000 storage unit to ESCON-attached System z and S/390 hosts.

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

- A maximum of 32 host ports for DS8100 Models 921 and 931 and a maximum of 64 host ports for DS8300 Models 922, 932, 9A2, and 9B2
- 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 System z and S/390 hosts:

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

For the most current information on supported hosts, operating systems, adapters and switches, see the DS8000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/> or the System Storage Interoperation Center (SSIC) at <http://www.ibm.com/systems/support/storage/config/ssic/>.

FICON-attached System z and S/390 hosts overview

This section describes how you can attach the DS8000 storage unit to FICON-attached System z and S/390 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 FICON upper-layer protocol. For FICON, the fibre-channel port supports connections to a maximum of 509 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 adapter ports on DS8100 Models 921 and 931, and a maximum of 128 host adapter ports on DS8300 Models 922, 932, 9A2, and 9B2
- A maximum of 509 logins per fibre-channel port
- A maximum of 8192 logins per storage unit
- A maximum of 1280 logical paths on each fibre-channel port
- Access to all 255 control-unit images (8000 CKD devices) over each FICON port
- A maximum of 512 logical paths per control unit image.

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. With this method, you can expose 64 control-unit images (16 384 devices) to each host channel.

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

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

For the most current information on supported hosts, operating systems, adapters and switches, see the DS8000 *Interoperability Matrix* at <http://www.ibm.com/>

servers/storage/disk/ds8000/ or the System Storage Interoperation Center (SSIC) at <http://www.ibm.com/systems/support/storage/config/ssic/>.

Processor memory

The DS8100 Models 921 and 931 offer up to 128 GB of processor memory. The DS8300 Models 922, 932, 9A2, and 9B2 offer up to 256 GB of processor memory. The nonvolatile storage (NVS) scales to the selected processor memory size, which can also help optimize performance.

Overview of DS8000 power controls

The DS8000 series provides power controls on the model racks. Power controls can be configured by an IBM service support representative. Power controls can also be accessed through the Management Console.

The DS8000 models have these following manual power controls in the form of physical switches on the racks:

- **Local/remote switch**

(Available on base models) The local/remote switch setting determines whether you will use local or remote power controls. When you set the switch to local, the local power on/local force power off switch controls power in the storage unit. You can access this switch by opening the rear cover of the unit. When the local/remote switch is set to remote, the power for the storage unit is controlled by remote power control settings that are entered in the DS Storage Manager.

Planning requirements: None.

- **Local power on/local force power off switch**

(Available on base models) The local power on/local force power off switch initiates a storage unit power on sequence or a storage unit force power off sequence. This switch is applicable only when the local/remote switch is set to local. You can access this switch by opening the rear cover of the unit. **Planning requirements:** None.

Note: To activate this switch, press and hold it for 8 seconds.

- **Unit emergency power off switch**

(Available on all models) *Use this switch only in extreme emergencies. Using this switch often results in data loss.* If activated, the unit emergency power off (UEPO) switch causes the individual model rack to immediately drop all power, including any power that is provided by the battery system. When active, this switch overrides all other power controls for the specific rack. This switch is located on the front of each model rack. **Planning requirements:** None.

The following power controls can be configured by an IBM service support representative. You can also use the following power controls through the DS Storage Manager (running on the Management Console):

- **Local power control mode**

(Visible in the DS Storage Manager) You cannot change this setting in the DS Storage Manager. This mode is enabled when the local/remote switch on the storage unit is in the local position. When this setting is used, the local power on/local force power off switch on the storage unit controls the power. **Planning requirements:** None.

- **Remote power control mode**

(Visible in the DS Storage Manager) If you select the **Remote** power control mode, you choose one of the following remote mode options. **Planning**

requirements: If you choose the **Remote zSeries Power Control** options, you must have the remote zSeries power control feature. There are no requirements for the other options.

- **Remote Management Console, Manual:** Your use of the DS Storage Manager power on/off page controls when the unit powers on and off.
- **Remote Management Console, Scheduled:** A schedule, which you set up, controls when the unit powers on and off.
- **Remote Management Console, Auto:** This setting applies only in situations in which input power is lost. In those situations, the unit powers on as soon as external power becomes available again.
- **Remote Auto/Scheduled:** A schedule, which you set up, controls when the unit powers on and off. A power on sequence is also initiated if the unit was powered off due to an external power loss during the time that the units are scheduled to be on and external power becomes available again.
- **Remote zSeries Power Control:** One or more attached System z or S/390 hosts can control the power on and power off sequences.

Powering a storage unit on and off (real-time only)

Complete this task to power on or power off a storage unit, bringing it online or taking it offline, respectively.

1. In the navigation, select **Real-time manager** → **Manage hardware** → **Storage units**. On the Storage units — Main Page, select the storage unit that you want to power on or off.
2. In the **Select Action** drop-down list, select **Power Control**. The Control mode page is displayed.
3. To enable the power on and off button for the following steps, set the **Current power control mode** to Remote and set the **Set remote power control mode** to Remote management console manual.
4. Click **Apply** and then click on the **Power on/off** option in the left navigation pane. The Power on/off page is displayed.
5. The button label depends on the current state of the storage unit. If the storage unit is currently off, this button has a label of “Power on”. If the storage unit is currently on, this button has a label of “Power off”. If you open this page while the storage unit is in the process of powering on or powering off, this button is disabled.
6. To change the state, press the power button. A confirmation message is displayed that asks you if you want to proceed. If you choose to continue, the power button is disabled and displays the label that is the opposite state of the box until the storage unit is completely powered on or off. A progress bar is displayed by the button to indicate that the process is working.
7. Once the process is complete, click **OK** to close the page.

Unit emergency power off (UEPO) switch

The UEPO switch is the red switch located on the rack operator panel. It must be used with extreme caution and only when the power must be dropped immediately for a safety concern.

The UEPO switch immediately shuts off all power to the machine. So, when the UEPO switch is turned off, all data not written to disk is lost and not recoverable.

Note: Using this switch often results in data loss.

Rack operator panel

The rack operator panel contains 3 light-emitting diodes (LEDs) and the unit emergency power off (UEPO) switch.

The LEDs reflect various conditions of operation depending on the function they represent. When you observe that their standard reflection is different, it can mean that there is a need for service.

Use the UEPO switch with extreme caution and only as a last resort.

The following graphic represents the rack operator panel that is found on a Model DS8000 (machine type 2107/242x).




Figure 6. Rack operator panel

The following identifies the LED indicators that are referenced in the graphic:

1. Line Cord 1 LED (for power system 1)
2. Line Cord 2 LED (for power system 2)
3. Attention LED

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 implementation

RAID implementation improves data storage reliability and performance.

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.

RAID implementation enables fault-tolerant data storage by storing the data in different places on multiple disk drive modules (DDMs). By placing data on multiple disks, input/output operations can overlap in a balanced way to improve the basic reliability and performance of the attached storage devices.

Physical capacity can be configured as RAID 5, RAID 6 (only on the DS8000 series), RAID 10, or a combination of RAID 5 and RAID 10. RAID 5 can offer excellent performance for most applications, while RAID 10 can offer better performance for selected applications, in particular, high random write content applications in the open systems environment. RAID 6 increases data protection by adding an extra layer of parity over the RAID 5 implementation.

You can reconfigure RAID 5 disk groups as RAID 10 disk groups or vice versa.

RAID 5 overview

RAID 5 is a method of spreading volume data across multiple disk drives. The DS8000 and DS6000 series supports RAID 5 arrays.

RAID 5 increases performance by supporting concurrent accesses to the multiple DDMs within each logical volume. Data protection is provided by parity, which is stored throughout the drives in the array. If a drive fails, the data on that drive can be restored using all the other drives in the array along with the parity bits that were created when the data was stored.

RAID 6 overview

RAID 6 is a method of increasing the data protection of arrays with volume data spread across multiple disk drives. The DS8000 series supports RAID 6 arrays.

RAID 6 increases data protection by adding an extra layer of parity over the RAID 5 implementation. By adding this protection, RAID 6 can restore data from an array with up to two failed drives. The calculation and storage of extra parity slightly reduces the capacity and performance compared to a RAID 5 array. RAID 6 is particularly suitable for storage using archive class DDMs.

RAID 10 overview

RAID 10 provides high availability by combining features of RAID 0 and RAID 1. The DS8000 and DS6000 series supports RAID 10 arrays.

RAID 0 increases performance by striping volume data across multiple disk drives. RAID 1 provides disk mirroring which duplicates data between two disk drives. By combining the features of RAID 0 and RAID 1, RAID 10 provides a second optimization for fault tolerance.

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. Access to data is preserved if one disk in each mirrored pair remains available. 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. RAID 10 arrays across loops provides mirroring across two loops, which prevents loss of the array during loop failure. RAID 5 and RAID 6 arrays across loops provide a higher throughput to the array because both loops on the device adapter are used in parallel. The main difference between RAID 6 and the other supported arrays is the additional parity DDM allocated to RAID 6.

Storage system logical partitions

The DS8300 Models 9A2 and 9B2 exploit logical partition (LPAR) technology, allowing you to run two separate storage server images.

Each storage image has access to the following components:

- 50 percent of the processors is the default. With variable LPAR implementation on DS8300 Model 9B2, the split can be 25/75 percent or 75/25 percent.
- 50 percent of the processors is the default. With variable LPAR implementation on DS8300 Model 9B2, the split can be 25/75 percent or 75/25 percent.
- 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).

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. Where separate storage subsystems were previously required, this capability enables storage consolidation, which increases management efficiency and cost effectiveness.

Typically, the different machine code levels of a common release are compatible. However, certain levels of licensed machine code might not be compatible if the later licensed machine code level contains substantial enhancements. If your levels of licensed machine code are not compatible, you can contact IBM support to request instructions for installing the two different licensed machine code levels. In addition, if each logical partition runs a different level of machine code, you can contact IBM support to request a code level upgrade on either partition. Your requests will be evaluated and approved based on your particular configuration.

IBM System Storage Resiliency Family

IBM System Storage Resiliency Family is a set of products and features that can help you implement storage solutions to keep your business running 24 hours a day, 7 days a week. This set of products provides disaster recovery, data migration, and data duplication functions.

As a part of the IBM System Storage Resiliency Family, the DS8000 series supports Copy Service functions that contribute to the protection of your data. These functions are also supported on the IBM TotalStorage Enterprise Storage Server.

The following IBM System Storage Resiliency Family functions are available as optional features:

- Point-in-time copy, which includes IBM System Storage FlashCopy and Space-Efficient FlashCopy
The FlashCopy function enables you to make point-in-time, full volume copies of data, so that the copies are immediately available for read or write access. In System z environments, you can also use the FlashCopy function to perform data set level copies of your data.
- Remote mirror and copy, which includes the following functions:
 - IBM System Storage Metro Mirror (previously known as Synchronous PPRC)
Metro Mirror provides real-time mirroring of logical volumes between two DS8000 storage units 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 done.
 - IBM System Storage Global Copy (previously known as PPRC Extended Distance)
Global Copy is a nonsynchronous long-distance copy function where incremental updates are sent from the local to the remote site on a periodic basis.
 - IBM System Storage Global Mirror (previously known as Asynchronous PPRC)
Global Mirror is a long-distance remote copy function across two sites using asynchronous technology. Global Mirror processing is designed to provide 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.
 - IBM System Storage Metro/Global Mirror (a combination of Metro Mirror and Global Mirror)
Metro/Global Mirror is a three-site remote copy solution, which uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.
- Remote mirror and copy for System z environments, which includes z/OS Global Mirror

The point-in-time and remote mirror and copy features are supported across various IBM server environments such as System i™, System p™, and System z, as well as servers from Sun and Hewlett-Packard.

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

Point-in-time copy (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. In System z environments, you can also use the FlashCopy function to perform data set level copies of your data. You can use the copy with standard backup tools that are available in your environment to create backup copies on tape.

FlashCopy is an optional function. To use it, you must purchase one of the point-in-time 2244/239x function authorization features.

The FlashCopy function 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.

One of the main benefits of the FlashCopy function is that the point-in-time copy is immediately available for creating a backup of production data. The target volume is available for read and write processing so it can be used for testing or backup purposes. Data is physically copied from the source volume to the target volume using a background process. (A FlashCopy operation without a background copy is also possible, which allows only data that is modified on the source to be copied 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 storage units

The FlashCopy function 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 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 (up to 12) 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.

Resynchronizing FlashCopy volume pairs

Provides the ability to update an initial point-in-time copy of a source volume without having to recopy your entire volume.

Reverse restore

Reverses the FlashCopy relationship and copies data from the target volume to the source volume.

Remote mirror and copy

The remote mirror and copy feature is a flexible data mirroring technology that allows replication between a source volume and a target volume on one or two disk storage units. You can also issue remote mirror and copy operations to a group of source volumes on one logical subsystem (LSS) and a group of target volumes on another LSS. (An LSS is a logical grouping of up to 256 logical volumes for which the volumes must have the same disk format, either count key data or fixed block.)

Remote mirror and copy is an optional feature that provides data backup and disaster recovery. To use it, you must purchase at least one of the remote mirror and copy 2244/239x function authorization features. DS8000 storage units can participate in remote mirror and copy solutions with the DS6000 series, ESS Model 750, and ESS Model 800.

The remote mirror and copy feature provides synchronous (Metro Mirror) and asynchronous (Global Copy) data mirroring. The main difference is that the Global Copy feature can operate at very long distances, even continental distances, with minimal impact on applications. Distance is limited only by the network and channel extenders technology capabilities. The maximum supported distance for Metro Mirror is 300 km.

With Metro Mirror, application write performance is dependent on the available bandwidth. Global Copy allows you to better use your available bandwidth capacity, therefore allowing you to include more of your data to be protected.

The enhancement to Global Copy is Global Mirror, which uses Global Copy and the benefits of FlashCopy to form consistency groups. (A consistency group is a set of volumes that contain consistent and current data to provide a true data backup at a remote site.) Global Mirror uses a master storage unit (along with optional subordinate storage units) to internally, without external automation software, manage data consistency across volumes using consistency groups.

Consistency groups can also be created using the freeze and run functions of Metro Mirror. The freeze and run functions, when used with external automation software, provide data consistency for multiple Metro Mirror volume pairs.

The following sections describe the remote mirror and copy functions.

Synchronous mirroring (Metro Mirror)

Provides real-time mirroring of logical volumes (a source and a target)

between two storage units that can be located up to 300 km from each other. 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.

Metro Mirror is a synchronous copy feature where write operations are completed on both copies (local and remote site) before they are considered to be complete. Synchronous mirroring means that a storage server constantly updates a secondary copy of a volume to match changes made to a source volume.

The advantage of synchronous mirroring is that there is minimal host impact for performing the copy. The disadvantage is that since the copy operation is synchronous, there can be an impact to application performance because the application I/O operation is not acknowledged as complete until the write to the target volume is also complete. The longer the distance between primary and secondary storage units, the greater this impact to application I/O, and therefore, application performance.

Asynchronous mirroring (Global Copy)

Copies data nonsynchronously and over longer distances than is possible with the Metro Mirror feature. 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.

The updates are tracked and periodically copied to the target volumes. As a consequence, there is no guarantee that data is transferred in the same sequence that was applied to the source volume. To get a consistent copy of your data at your remote site, you must periodically switch from Global Copy to Metro Mirror mode, then either stop the application I/O or freeze data to the source volumes using a manual process with freeze and run commands. The freeze and run functions can be used with external automation software such as Geographically Dispersed Parallel Sysplex™ (GDPS®), which is available for System z environments, to ensure data consistency to multiple Metro Mirror volume pairs in a specified logical subsystem.

Common options for Metro Mirror and Global Copy include the following modes:

Suspend and resume

If you schedule a planned outage to perform maintenance at your remote site, you can suspend Metro Mirror or Global Copy processing on specific volume pairs during the duration of the outage. During this time, data is no longer copied to the target volumes. Because the primary storage unit keeps track of all changed data on the source volume, you can resume operations at a later time to synchronize the data between the volumes.

Copy out-of-synchronous data

You can specify that only data that was updated on the source volume while the volume pair was suspended be copied to its associated target volume.

Copy an entire volume or not copy the volume

You can copy an entire source volume to its associated target volume to guarantee that the source and target volume contain the

same data. When you establish volume pairs and elect not to copy a volume, a relationship is established between the volumes but no data is sent from the source volume to the target volume. In this case, it is assumed that the volumes contain exactly the same data and are consistent, so copying the entire volume is not necessary or required. Only new updates are copied from the source to target volumes.

Global Mirror

Provides a long-distance remote copy across two sites using asynchronous technology. Global Mirror processing is most often associated with disaster recovery or disaster recovery testing. However, it can also be used for everyday processing and data migration.

The Global Mirror function mirrors data between volume pairs of two storage units over greater distances without affecting overall performance. It also provides application-consistent data at a recovery (or remote) site in case of a disaster at the local site. By creating a set of remote volumes every few seconds, the data at the remote site is maintained to be a point-in-time consistent copy of the data at the local site.

Global Mirror operations periodically invoke point-in-time FlashCopy operations at the recovery site, at regular intervals, without disrupting the I/O to the source volume, thus giving a continuous, near up-to-date data backup. 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. (A session contains a group of source volumes that are mirrored asynchronously to provide a consistent copy of data at the remote site. Sessions are associated with Global Mirror relationships and are defined with an identifier [session ID] that is unique across the enterprise. The ID identifies the group of volumes in a session that are related and that can participate in the Global Mirror consistency group.)

Metro/Global Mirror

Provides a three-site, long distance disaster recovery replication that combines Metro Mirror with Global Mirror replication for both System z and open systems data. Metro/Global Mirror uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

In a three-site, Metro/Global Mirror, should an outage occur, a backup site is maintained regardless of which one of the sites is lost. Suppose an outage occurs at the local site, Global Mirror continues to mirror updates between the intermediate and remote sites, maintaining the recovery capability at the remote site. If an outage occurs at the intermediate site, data at the local storage unit is not affected. If an outage occurs at the remote site, data at the local and intermediate sites is not affected. Applications continue to run normally in either case.

With the incremental resynchronization function enabled on a Metro/Global Mirror configuration, should the intermediate site be lost, the local and remote sites can be connected, and only a subset of changed data is copied between the volumes at the two sites. This reduces the amount of data that needs to be copied from the local site to the remote site and the time it takes to do the copy.

z/OS Global Mirror

In the event of workload peaks, which might temporarily overload the bandwidth of the Global Mirror configuration, the enhanced z/OS Global Mirror function initiates a Global Mirror suspension that preserves primary site application performance. If you are installing new high-performance z/OS Global Mirror primary storage subsystems, this function provides improved capacity and application performance during heavy write activity. This enhancement can also allow Global Mirror to be configured to tolerate longer periods of communication loss with the primary storage subsystems. This enables the Global Mirror to stay active despite transient channel path recovery events. In addition, this enhancement can provide fail-safe protection against application system impact related to unexpected data mover system events.

The z/OS Global Mirror function is an optional function. To use it, you must purchase the remote mirror for z/OS 2244/239x function authorization feature.

z/OS Global Mirror Multiple Reader (enhanced readers)

z/OS Global Mirror Multiple Reader provides multiple Storage Device Manager readers allowing improved throughput for remote mirroring configurations in System z environments. z/OS Global Mirror Multiple Reader helps maintain constant data consistency between mirrored sites and promotes efficient recovery. This function is supported on the DS8000 series running in a System z environment with version 1.7 or later at no additional charge.

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/Global Mirror	Three-site, long distance disaster recovery replication	A backup site is maintained regardless of which one of the sites is lost.	Recovery point objective (RPO) might grow if bandwidth capability is exceeded.
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.

Table 2. Comparison of licensed functions (continued)

Licensed function	Description	Advantages	Considerations
Global Mirror	Asynchronous copy	Nearly unlimited distance, scalable, and low 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.
z/OS Global Mirror	Asynchronous copy controlled by System z 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 and IBM HyperPAV

The use of the parallel access volumes (PAV) feature enables a single System z or S/390 host 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 PAV, 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.

For System z or S/390 hosts, the storage unit supports concurrent or parallel data transfer operations to or from the same volume from the same system or system image. 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.)

You can enhance PAV by adding the IBM HyperPAV feature which allows you to use alias addresses to access any base on the same control unit image per I/O base.

CKD extended address volumes - 3390 Model A

The DS8000 supports CKD extended address volumes and provides increased storage capacity with 3390 Model A.

3390 Model A is intended as replacement for 3390 Model 3 and 3390 Model 9. Volumes can be converted between 3390 Model 3, 3390 Model 9, and 3390 Model A if the volume sizes are within 3390 Model 3 and 3390 Model 9 limits. 3390 Model 3, 3390 Model 9, and 3390 Model A volumes can be increased in size to their

maximum supported sizes. DS8000 and z/OS limit CKD extended address volumes sizes. A 3390 Model A volume has the capacity of 1 - 262668 cylinders (approximately 223.2 GB maximum).

z/VM and z/OS software are required to use extended address volumes. Volume creation and volume expansion is limited to 3390 Model 9 sizes if earlier software has a path group on the DS8000. Volumes larger than 3390 Model 9 do not come online if earlier z/OS software is in use.

DS8000 series limitations

The following list describes known limitations for the DS8000 series.

- The 65 520 cylinder 3390 volume is not supported with z/OS Global Mirror.
- The IBM FlashCopy SE function is not supported in DS8000 machines that have a 3rd and 4th expansion unit attached.
- For the Dynamic Volume Expansion function, volumes cannot be in Copy Services relationships (point-in-time copy, FlashCopy SE, Metro Mirror, Global Mirror, Metro/Global Mirror, and z/OS Global Mirror) during expansion.
- The amount of physical capacity within a 2107/242x system that can be logically configured for use will be enforced by the 2107/242x licensed machine code (LMC) to maintain compliance with the extent of IBM authorization established for licensed functions activated on the machine. The 2107/242x LMC will not allow the logical configuration of physical capacity beyond the extent of IBM authorization (except when activating Standby CoD capacity).
- The 500 GB FATA 7 200 rpm disk drives support various fixed content, data archival, reference data, and near-line applications that require large amounts of data at low cost. They are not intended for use in applications that require drive duty cycles greater than 20%. To intermix the FATA disk drive sets and fibre-channel disk drive sets is not supported within a disk enclosure or at creation. However, intermixing on the same DA loop of enclosures with different drive types is supported if the drives within each enclosure are the same.
- The deactivation of an activated licensed function, a lateral change, and the reduction in the license scope are all disruptive activities and require a machine IML (Models 921, 922, 931, and 932) or reboot of the affected image (Model 9A2 and 9B2). 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:
 - Addition of the Earthquake Resistance Kit feature 1906.
 - Addition of the 150x POWER5+™ features.
 - All model conversions. In addition, data is not preserved during the model conversion unless the conversion is 921/931.
 - Field attachment of the first expansion (Model 92E) to a base (Model 922 or 932).
 - Field attachment of the first expansion (Model 9AE) to a base (Model 9A2 or 9B2).
 - Removal of an expansion model from the base model. Data is not preserved during this activity.
- Some DS8000 series 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:

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 a System z or UNIX-based system? You might use IBM System Storage Remote Mirror and Copy functions such as Metro Mirror and Global Mirror or some variation of a logical volume manager.
 - Is it System z? 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.

Logical configuration overview

Before you configure your DS8000, it is important to understand IBM terminology for storage concepts and the storage hierarchy.

In the storage hierarchy, you begin with a disk. Logical groupings of eight disks form an array site. Logical groupings of one array site form an array. After you define your array storage type as CKD or fixed block, you can create a rank. A rank is divided into a number of fixed-size extents. If you work with an open-systems host, an extent is 1 GB. If you work in an IBM System z environment, an extent is the size of an IBM 3390 Mod 1 disk drive.


After you create ranks, your physical storage can be considered virtualized. Virtualization dissociates your physical storage configuration from your logical configuration, so that volume sizes are no longer constrained by the physical size of your arrays.

The available space on each rank is divided into extents. The extents are the building blocks of the logical volumes. An extent is striped across all disks of an array.

Extents of the same storage type are grouped together to form an extent pool. Multiple extent pools can create storage classes that provide greater flexibility in storage allocation through a combination of RAID types, DDM size, DDM speed, and DDM technology. This allows a differentiation of logical volumes by assigning them to the appropriate extent pool for the desired characteristics. Different extent sizes for the same device type (for example, count-key-data or fixed block) can be supported on the same storage unit, but these different extent types must be in different extent pools.

A logical volume is composed of one or more extents. A volume group specifies a set of logical volumes. By identifying different volume groups for different uses or functions (for example, SCSI target, FICON/ESCON control unit, remote mirror and copy secondary volumes, FlashCopy targets, and Copy Services), access to the set of logical volumes that are identified by the volume group can be controlled. Volume groups map hosts to volumes. Figure 7 on page 34 shows a graphic representation of the logical configuration sequence.

When volumes are created, you must initialize logical tracks from the host before the host is allowed read and write access to the logical tracks on the volumes. An internal volume initialization process allows quicker access to logical volumes that are used as host volumes and source volumes in Copy Services relationships, such as FlashCopy or Remote Mirror and Copy relationships. Volumes are not available as target volumes for Copy Services relationships until the volumes have been fully initialized. After volume initialization is complete and successful, you can use the target volumes to establish Copy Services relationships. Initialization requests can fail if the initialization is performed on volumes that are participating in existing Copy Services relationships.



f2d00137

Figure 7. Logical configuration sequence

The storage management software can be used in real-time or simulated modes. When you are connected to storage devices over your network, you can use the Real-time Manager to manage your hardware or configure your storage. When you want to work offline, you can use the Simulated Manager to create or import a configuration. This lets you export a configuration that you have already created to a system when it reaches your site.

Chapter 2. Installing the simulated DS8000 Storage Manager interface

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

You can install the simulated 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	1001 MB
Memory	1024 MB Real Memory
Processor	Pentium 4 Processor at 1.4 GHz

Supported browsers:

- Internet Explorer 6.x and 7.x
- Mozilla Firefox 1.5 and 2.0

Notes:

1. To prevent the content of your browser window from being replaced by other windows being launched, check your internet options to ensure that the option to reuse windows for shortcuts is disabled.
2. 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. Animations are supported by default in Mozilla Firefox. Use the following instructions for Internet Explorer:
 - a. From the **Tools** menu, select **Internet Options**.
 - b. Select the **Advanced** tab and scroll down to the **Multimedia** section.
 - c. Check **Play animation in web pages**.

Supported operating systems:

- Windows 2000 Professional, Server, Advanced Server
- Windows 2003 Server, Advanced Server
- Windows XP Professional
- Red Hat Linux AS 3.0 and 4.0

Notes:

1. The simulated DS8000 Storage Manager is not supported on any Windows 64-bit operating system.
2. If this installation is being done on top of an existing installation, the installer will automatically remove WebSphere Application Server 5 and install WebSphere Application Server 6. The WebSphere Application Server kit contains pre-upgrade and post-upgrade tools that you can use for problem determination if necessary. Use the commands *WASPostUpgrade* or *WASPreUpgrade* to identify any errors that might

have occurred during the upgrade. You can check online support for error definitions. Additional information can be found in the logs of the local directories.

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 simulated DS8000 Storage Manager, the following results occur:

- Activation of the IBM System Storage DS Storage Manager server and the IBM System Storage 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 System Storage DS8000 Storage Manager application, which includes the simulated storage manager component. It is designed to help you create and manage the physical and logical configurations of your storage complexes, storage units, and storage images.

This component (installed using the IBM System Storage DS8000 Storage Manager installation CD) provides the ability to create 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.

Attention: The simulated DS Storage Manager is designed to assist you in creating an initial configuration for your system. Do not attempt to upload a configuration file if your system has already been configured.

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

Complete this task to install the simulated IBM System Storage DS Storage Manager in your Windows environment using the graphical mode. You are 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 simulated IBM System Storage DS Storage Manager, verify that the prerequisite software and hardware are installed on your system. The installation program checks for prerequisites and cancels the installation if any prerequisites are missing. If you have the correct level of Java code, but Java errors appear during installation, request a new installation CD from your IBM representative.

Notes:

1. The name of the user who logs on and performs the installation must not contain any non-ASCII or special characters. It could cause the installation to fail.
2. The DS Storage Manager is not supported on any Windows 64-bit operating system.

To install the DS Storage Manager on the Windows operating system using graphical mode, perform the following steps:

1. Log on as a user with administrator authority.
2. Insert the IBM System Storage DS Storage Manager CD into the CD-ROM drive. The IBM System Storage 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:

- Type *LaunchPad* at the command prompt and press **Enter**. The LaunchPad window is displayed.
- 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.

The following options are available in the Launchpad window:

DS8000 Storage Manager overview

Provides information about the IBM System Storage DS Storage Manager software.

Readme file (recommended selection)

Provides updated 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 System Storage DS Storage Manager installation program.

Post installation tasks

Provides information about setting up the IBM System Storage DS Storage Manager.

Exit

Exits the IBM System Storage DS Storage Manager LaunchPad program.

3. Click the **Readme file** selection on the LaunchPad to check for information that might supersede the information in this guide.
4. 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 System Storage 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.

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

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.

6. 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.
7. Choose the directory on the Destination Directory window where you want to install the application by using one of the following options. Clicking **Next** causes the server installation checking window (a window for each server) to display if the DS Storage Manager Server or the DS Network Interface Server is installed on your system. If the two servers are not installed on your system, then the Server Parameters window is displayed.
 - Click the **Next** button to accept the default (recommended) directory that is shown in the window.
 - Type a fully qualified directory path into the Directory Name path to designate your own directory. Then click the **Next** button.
 - 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.
8. 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 9 to proceed with the Server Parameters window. For a typical installation, default server parameters can be used.

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 15 on page 41.

9. 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.

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

- **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 the 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*.

- **HTTP Port** — The port number that the DS Storage Manager Server uses. It is recommended that the port numbers specified in this section not be changed from the defaults. Doing so, might compromise remote support. 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`.
- **HTTPS Port** — The port that the DS Storage Manager Server uses for secure HTTP transport (HTTPS).
- **Bootstrap/RMI Port** — The port that the DS Storage Manager Server uses for the bootstrap port.
- **SOAP Port** — The port that the DS Storage Manager Server uses for Simple Object Access Protocol (SOAP).
- **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.

10. 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 DS Storage Manager installation.
11. Follow the instructions on the SSL Configuration window.

- a. 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 DS 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. If you are creating new certificates, create a password for each file. The password must be a minimum of six characters and it cannot contain any blanks. Allowable characters are letters A - Z and a - z, digits 1 - 9, and special characters - and _ (hyphen and underscore). Supply the existing password if you are using an existing SSL certificate.

Record the location and passwords for the SSL certificates so that you can refer to them in the future if necessary.

12. 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 keep existing application server certificates if they exist on the location. The default is **YES** which keeps the existing application server certificates.

13. 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.

14. 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.

Note: The total size required for installation might vary, depending on your installation environment.

15. 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:
 - 1) The Embedded IBM WebSphere Application Server - Express server (part of the DS Storage Manager Server)
 - 2) The DS Network Interface Server
 - 3) The WS Help System (part of the DS Storage Manager Server), 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 must 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 of the component files that have been copied up to the point that the **Cancel** button was clicked are uninstalled and the installation process is cancelled.

When this part of the installation is completed, if no system restart is required, 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 with the result of the installation process (successful or unsuccessful).

If a system restart is required, the Installer Finish window is displayed, which asks you to restart the system. Click **Next** on the Finish window. The Installer Reboot window is displayed. Choose the default selection **Yes, restart my computer** on the Reboot window, and then click **Finish**. Both servers, first the DS Network Interface Server and then the DS Storage Manager Server, start automatically after the restart.

16. Click **Finish** to exit the installation process. When the installation process exits, a file (postinstallation.txt) is displayed, containing the post installation tasks. Use these instructions to complete the setup tasks.

Note: If the Finish window indicates that the installation fails, you must exit the installation process and 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 System Storage DS Storage Manager is installed (for example, `c:\Program Files\IBM\DS8000 Storage Manager`).

Errors fall into two categories: system problems 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 DS Storage Manager using either the interactive or unattended (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 unattended (silent) mode of installation. Navigate to the Server Parameters window and insert the correct values. Then finish the installation process.
17. Start the DS Storage Manager. See “Starting the DS Storage Manager on a Windows operating system” on page 81.
 18. 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 post installation tasks.
 19. Exit the LaunchPad program by clicking **Exit** on the LaunchPad window.

Installing the IBM System Storage DS Storage Manager on the Windows operating system using unattended (silent) mode

Complete this task to install the simulated IBM System Storage DS Storage Manager in your Windows environment using the unattended (silent) mode. When you are installing on a DS8000 storage unit, you are 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 System Storage 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. If you have the correct level of Java code, but Java errors appear during installation, request a new installation CD from your IBM representative.

Notes:

1. The name of the user who logs on and performs the installation must not contain any non-ASCII or special characters. It could cause the installation to fail.
2. The DS Storage Manager is not supported on any Windows 64-bit operating system.

Using 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 System Storage DS Storage Manager CD.

Perform the following steps to install the IBM System Storage DS Storage Manager in your Windows environment using the unattended (silent) mode:

1. Log on to your Windows system as an administrator.
2. Insert the IBM System Storage 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 password must be a minimum of six characters and it cannot contain any blanks.
 - Allowable characters are letters A - Z and a - z, digits 1 - 9, and special characters - and _ (hyphen and underscore).
 - Use the following options within the responsefile.txt file to set these options:
 - For the keystore file use:

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

- For the truststore file 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 unattended (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. If required, perform a system restart to complete the installation. After the system has restarted, the two servers (DS Storage Manager server and DS Network Interface Server) are automatically started by the system.
9. Verify that the DS Storage Manager server and IBM System Storage DS Network server are operating. From the Windows Control Panel, select

Administrative Tools, and then **Services**. Both servers must show a "Started" status. You cannot start the DS Storage Manager until both of the servers are started.

10. 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 System Storage DS Storage Manager is installed.
11. Start the IBM System Storage DS Storage Manager, if no errors are evident.
12. Perform the post installation tasks when the DS Storage Manager has been installed successfully. You can access the instructions for the post installation 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 Template
#
# Wizard name: Setup
# Wizard source: setup.jar
# Created on: Mon Jul 30 17:58:51 EEST 2007
# Created by: InstallShield Options File Generator
#
# This file can be used to create an options file (i.e., response file) for the
# wizard "Setup". Options files are used with "-options" on the command line to
# modify wizard settings.
#
# The settings that can be specified for the wizard are listed below. To use
# this template, follow these steps:
#
# 1. Enable a setting below by removing leading '###' characters from the
# line (search for '###' to find settings you can change).
#
# 2. Specify a value for a setting by replacing the characters '<value>'.
# Read each setting's documentation for information on how to specify its
# value.
#
# 3. Save the changes to the file.
#
# 4. To use the options file with the wizard, specify -options <file-name>
# as a command line argument to the wizard, where <file-name> is the name
# of this options file.
#
#####
```

```
#####
#
# Has the license been accepted
#
# The license must be accepted before installation so this value must be true
# for the install to be successful. Example: -G licenseAccepted=true
#
### -G licenseAccepted=<value></value>
```

```
#####
#
# IBM System Storage 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. For example, to install the product to C:\Program Files\My
# Product, use
#
# -P installLocation="C:\Program Files\My Product"
#
### -P installLocation=<value>
```

```
#####  
#  
# User Input Field - hostname  
#  
# The fully-qualified host name of the machine where the DS Storage Manager  
# Server will be installed.  
#  
### -W wasExpressConfig.hostname="<value>"</value>
```

```
#####  
#  
# User Input Field - httpPort  
#  
# The port number that the HTTP transport in the application server will use.  
# The HTTP transport is a request queue between the application server and the  
# HTTP server (Web server). This value must not conflict with existing port  
# assignments on the system.  
#  
### -W wasExpressConfig.httpPort="<value>"
```

```
#####  
#  
# 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="<value>"
```

```
#####  
#  
# 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="<value>"
```

```
#####  
#  
# User Input Field - helpPort  
#  
# The port that the help system (based on Eclipse technology) will use to  
# receive requests for help files. This value must not conflict with existing  
# port assignments on the system.  
#  
### -W wasExpressConfig.helpPort="<value>"
```

```
#####  
#  
# User Input Field - certificateOption  
#  
# SSL Certificates Files 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 wasSslConfiguration.certificateOption="generate"  
#  
### -W wasSslConfiguration.certificateOption="<value>"
```

```
#####  
#  
# User Input Field - keystoreFileName  
#  
# The absolute path of the keystore file.  
#  
### -W wasSslConfiguration.keystoreFileName="<value>"
```

```
#####  
#  
# User Input Field - keystorePassword  
#  
# The password for the keystore file.  
#  
### -W wasSslConfiguration.keystorePassword="<value>"</value>
```

```
#####  
#  
# User Input Field - confirmedKeystorePassword  
#  
# The password confirmation of the keystore file. On silent mode the password  
# confirmation is NOT necessary.  
#  
### -W wasSslConfiguration.confirmedKeystorePassword="<value>"
```

```
#####  
#  
# User Input Field - truststoreFileName  
#  
# The absolute path of the truststore file.  
#  
### -W wasSslConfiguration.truststoreFileName="<value>"</value>
```

```
#####  
#  
# User Input Field - truststorePassword  
#  
# The password for the truststore file.  
#  
### -W wasSslConfiguration.truststorePassword="<value>"
```

```
#####  
#  
# User Input Field - confirmedTruststorePassword  
#  
# The password confirmation of the truststore file. On silent mode the password  
# confirmation is NOT necessary.  
#  
### -W wasSslConfiguration.confirmedTruststorePassword="<value>"
```

```
#####  
#  
# 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=no
```

```
#####  
#  
# 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="<value>"
```

```
#####  
#  
# 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="<value>"</value>
```

```
#####  
#  
# 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="<value>"
```

```
#####  
#  
# User Input Field - organizationName  
#  
# The name of your organization.  
#  
  
### -W CertificateParams.organizationName="<value>"
```

```
#####  
#  
# 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="<value>"
```

```
#####  
#  
# 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="<value>"
```

```
#####  
#  
# 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="<value>"
```

```
#####  
#  
# 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="<value>"
```

```
#####  
#  
# User Input Field - validity  
#  
# The lifetime of the certificate in days.  
#  
  
### -W CertificateParams.validity="<value>"
```

Removing the DS Storage Manager from a Windows operating system

Complete this task to remove the IBM System Storage DS Storage Manager from your Windows operating system by using the Windows Add/Remove Programs facility.

1. Log on to your Windows operating system with administrator authority.
2. Navigate to your Windows Control Panel and open the Add/Remove program facility.
3. Scroll the list of currently installed programs and click the listing for IBM System Storage DS8000 Storage Manager.
4. Click the **Change/Remove** button, and the Welcome window for the Uninstaller is displayed.
5. Click **Next** to continue or click **Cancel** to exit the removal process. When you click **Next**, the Confirmation window is displayed that shows the directory from which the IBM System Storage DS Storage Manager GUI program will be removed.
6. 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 V6.1 - DS Storage Manager** and **IBM DS Network Interface Server** services to the **Manual** startup type.
 - b. Reboot the computer.
 - c. Proceed once again with the removal of the DS Storage Manager.

Another reason that files can be locked is because you changed the host name where 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.

7. Click **Finish** to complete the removal process and exit the GUI part of the uninstallation program.
8. Close the Add/Remove Programs window.

- Restart your system, if required (now or later), to complete the removal process.

If you want to remove all configuration settings before you reinstall the DS Storage Manager, you must delete the `dsnserver` and the `SMServer` subdirectories from the `Program Files\IBM` directory.

Removing the DS Storage Manager from a Windows operating system using the unattended (silent) mode

Complete this task to remove the IBM System Storage DS Storage Manager from your Windows operating system by using the unattended (silent) mode.

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

If you want to remove all configuration settings before you reinstall the DS Storage Manager, you must delete the `dsnserver` and the `SMServer` subdirectories from the `Program Files\IBM` directory.

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

Complete this task 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. It is recommended that you upgrade the level of code before you upgrade the DS Storage Manager.

Notes:

- The name of the user who logs on and performs the installation must not contain any non-ASCII or special characters. It could cause the installation to fail.
- The DS Storage Manager is not supported on any Windows 64-bit operating system.

To upgrade the DS Storage Manager on a Windows operating system using the graphical mode, perform the following steps:

- Log on as a user with administrator authority.
- If you are installing from a CD, insert the IBM System Storage DS Storage Manager CD into the CD-ROM drive. The IBM System Storage DS Storage Manager program starts within 15 - 30 seconds if you have autorun mode set on your system. The LaunchPad window opens.

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 opens.
- 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 that hides the extensions for unknown file types, use the LaunchPad file with the file type of MS-DOS Batch file.

The following options are available in the Launchpad window:

DS8000 Storage Manager overview

Provides information about the IBM System Storage 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 System Storage DS Storage Manager installation program.

Post installation tasks

Provides information about configuring the IBM System Storage DS Storage Manager.

Exit

Exits the IBM System Storage DS Storage Manager LaunchPad program.

3. If you are installing from a CD, click the **Readme file** selection on the LaunchPad to check for updated information.
4. 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 System Storage 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.

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

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

6. 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 opens.
7. 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 opens if the DS Storage Manager service is installed on your system.
8. 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 opens if the DS Network Interface service is installed on your system.
9. 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 opens. This window displays both the location where the product will be installed and the total size that is required for the installation.
10. Click **Next**. The SSL Configuration window is displayed. Enter the fully qualified name of the two server key files generated before or during the DS Storage Manager installation.
11. Follow the instructions on the SSL Configuration window:
 - a. Select either Generate the Self-signed Certificates during installation or select Use Existing SSL Certificates. The most likely selection when you are installing for the first time is Generate the Self-signed Certificates. The selection Use Existing SSL Certificates 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. Use the Browse button to provide the fully qualified path to the key file and trust file. The key file and trust file are identified by an extension of .jks.
 - c. If you are creating new certificates, create a password for each file. The password must be a minimum of six characters and it cannot contain any blanks. Allowable characters are letters A - Z and a - z, digits 1 - 9, and special characters - and _ (hyphen and underscore). Supply the existing password if you are using an existing SSL certificate. Record the location and passwords for the SSL certificates so that you can refer to them in the future if necessary.
12. Click **Next** to continue the installation process. If you 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. 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 keep existing application server certificates if they exist on the location. The default is YES which keeps the existing application server certificates.

13. 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.

14. Click **Next** to continue, or click **Cancel** to exit the installation.
15. Click the **Install** button on the Installation Confirmation window to begin the installation process. There are several progress windows that open. 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:
 - 1) The Embedded IBM WebSphere Application Server - Express server (part of the DS Storage Manager Server) is stopped first.
 - 2) The DS Network Interface Server is stopped next.
 - 3) The WS Help System (part of the DS Storage Manager Server) 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 opens. 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.

When this part of the installation is completed, if no system restart is required, 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 opens with the result of the installation process (successful or unsuccessful).

If a system restart is required, the Installer Finish window is displayed, which asks you to restart the system. Click **Next** on the Finish window. The Installer Reboot window is displayed. Choose the default selection **Yes, restart my computer** on the Reboot window, and then click **Finish**. Both servers, first the DS Network Interface Server and then the DS Storage Manager Server, start automatically after the restart.

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

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

- 17. 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 System Storage 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 unattended (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 unattended (silent) mode installation process. Navigate to the Server Parameters window and insert the correct values. Finish the installation process.
18. Complete the post installation 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 post installation tasks.
 19. End the LaunchPad program by clicking **Exit** on the LaunchPad window.

Upgrading the DS Storage Manager on the Windows operating system using unattended (silent) mode

Complete this task to upgrade the IBM System Storage DS Storage Manager in your Windows environment using the unattended (silent) mode.

It is recommended that you upgrade the level of code before upgrading the DS Storage Manager.

Notes:

1. The name of the user who logs on and performs the installation must not contain any non-ASCII or special characters. It could cause the installation to fail.
2. The DS Storage Manager is not supported on any Windows 64-bit operating system.

Using 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 System Storage DS Storage Manager CD.

Perform the following steps to upgrade the IBM System Storage DS Storage Manager in your Windows environment using the unattended (silent) mode:

1. Log on to your Windows system as an administrator.
2. Insert the IBM System Storage DS Storage Manager upgrade CD into the CD-ROM drive. If the interactive mode starts, click the **Exit** selection on the LaunchPad window to end the program.
3. Type the following command at the command prompt: `setup.exe -silent`. Press the **Enter** key to start the upgrade process in unattended (silent) mode.
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 DS Storage Manager is installed.
5. Start the IBM System Storage 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 post installation tasks from the `postinstallation.txt` file, in the `doc` directory on the product CD.

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

You can choose to install the simulated IBM System Storage DS8000 Storage Manager on Red Hat Linux AS 3.0 or 4.0 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 System Storage DS Storage Manager on a Linux operating system using the graphical mode

You can install the IBM System Storage 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 System Storage 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.

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

DS8000 Storage Manager overview

Provides information about the IBM System Storage 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 System Storage DS Storage Manager installation program.

Post installation tasks

Provides information about configuring the IBM System Storage DS Storage Manager.

Exit

Exits the IBM System Storage 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 System Storage 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 System Storage DS Storage Manager installation program displays.

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

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.
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 of the options causes the server installation checking window (a window for each server) to display if the DS Storage Manager Server or the DS Network Interface Server is installed on your system. If the two servers are not installed on your system, then 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 the two servers are already installed on your system. If your installation process has taken you through these two checking windows, go to Step 19 on page 62 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:
 - 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*.

- 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`
- HTTPS Port** — The port that DS Storage Manager Server uses for secure HTTP transport (HTTPS).
- Bootstrap/RMI Port** — The port that DS Storage Manager Server uses for bootstrap port.
- SOAP Port** — The port that DS Storage Manager Server uses for Simple Object Access Protocol (SOAP).
- 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.

- 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.
- Follow the instructions on the window.
 - 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.
 - 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. Allowable characters are letters A-Z and a-z, digits 1-9, and special characters - and _ (hyphen and underscore).

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 keep existing application server certificates if they exist on the location. The default is **YES**, which means that you want to keep the existing application server certificates.

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
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:
 - 1) The Embedded IBM WebSphere Application Server - Express server (part of the DS Storage Manager Server) is stopped first.
 - 2) The DS Network Interface Server is stopped next.
 - 3) The WS Help System (part of the DS Storage Manager Server) 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, if no system reboot is required, 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 with the result of the installation process (successful or unsuccessful).

If a system logoff is required, the Installer Finish window is displayed, which asks you to log off from the system. Click **Next** on the Finish window. The Installer Logoff window is displayed. Click **Finish** on the Logout window.

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.
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 System Storage 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 unattended (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 unattended (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 System Storage 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 System Storage DS8000 Storage Manager on the Linux operating system in unattended (silent) mode

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

Using 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 System Storage DS Storage Manager CD.

Perform the following steps to install the IBM System Storage DS Storage Manager in your Linux environment using the unattended (silent) 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 unattended (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 and it cannot contain any blanks.
 - Allowable characters are letters A-Z and a-z, digits 1-9, and special characters - and _ (hyphen and underscore).
 - 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>"
```
- 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 unattended (silent) mode.

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

<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 System Storage 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/  
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="/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 System Storage DS Storage Manager from a Linux operating system

You can remove the IBM System Storage 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. `<dest_path>` is the path where the IBM System Storage DS Storage Manager is installed. If you used the default path, it is `/opt/IBM/DS8000StorageManager`. The Uninstaller Welcome window is displayed.

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 System Storage DS Storage Manager GUI program is removed.
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.
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 unattended (silent) mode

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

1. Log on as root to the system where the DS8000 Storage Manager is installed.
2. Type `# <dest_path>_uninst\uninstall -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 System Storage 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 System Storage 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.

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

DS8000 Storage Manager overview

Provides information about the IBM System Storage 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 System Storage DS Storage Manager upgrade program.

Post installation tasks

Provides information about configuring the IBM System Storage DS Storage Manager.

Exit

Exits the IBM System Storage 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 System Storage 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 System Storage DS Storage Manager installation program displays.

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

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.
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.
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.
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.
14. Click **Next**. The SSL Configuration window is displayed. Enter the fully qualified name of the two server key files that are generated before or during the DS Storage Manager installation.
15. Follow the instructions on the SSL Configuration window:
 - a. Select either Generate the Self-signed Certificates during installation or select Use Existing SSL Certificates. The most likely selection when you are installing for the first time is Generate the Self-signed Certificates. The selection Use Existing SSL Certificates 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. Use the Browse button to provide the fully qualified path to the key file and trust file. The key file and trust file are identified by an extension of .jks.

- c. If you are creating new certificates, create a password for each file. The password must be a minimum of six characters and it cannot contain any blanks. Allowable characters are letters A - Z and a - z, digits 1 - 9, and special characters - and _ (hyphen and underscore). Supply the existing password if you are using an existing SSL certificate. Record the location and passwords for the SSL certificates so that you can refer to them in the future if necessary.
16. Click **Next** to continue the installation process. If you 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. 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 keep existing application server certificates if they exist on the location. The default is YES which keeps the existing application server certificates.
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 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.

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

19. 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:
 - 1) The Embedded IBM WebSphere Application Server - Express server (part of the DS Storage Manager Server) is stopped first.
 - 2) The DS Network Interface Server is stopped next.
 - 3) The WS Help System (part of the DS Storage Manager Server) 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.

If a system logoff is required, the Installer Finish window is displayed, which asks you to log off from the system. Click **Next** on the Finish window. The Installer Logoff window is displayed. Click **Finish** on the Logout window.

20. 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.

21. 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 System Storage 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 unattended (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 unattended (silent) mode upgrade. Navigate to the Console Server Parameters window and insert the correct values. Then finish the upgrade.
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 System Storage 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 System Storage DS8000 Storage Manager on the Linux operating system in unattended (silent) mode

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

Using 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 System Storage DS Storage Manager CD.

Perform the following steps to upgrade the IBM System Storage DS Storage Manager in your Linux environment using the unattended (silent) 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 unattended (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 different versions of Internet browsers.

The following list shows the Internet browser versions that support the use of the DS Storage Manager. Unless otherwise noted, these browsers work with the DS Storage Manager installed on the supported operating systems.

- Internet Explorer 6.x and 7.x
- Mozilla Firefox 1.5 and 2.0

Notes:

1. You must select the appropriate browser security settings to open DS Storage Manager in a browser. In Internet Explorer, on the Tools menu, click Internet Options. On the Security tab, select Internet and click Custom level. Under Miscellaneous, enable **Allow META REFRESH**. Under Scripting, enable **Active scripting**.
2. When running the DS8000 Element Manager from the IBM System Storage Productivity Center, with a combination of Internet Explorer 6 and SSL (HTTPS) mode, intermittent connection problems are seen. These problems can cause some usability inconveniences as well as commands not being run successfully.

If you access IBM TotalStorage Productivity Center through Java Web Start, TotalStorage Productivity Center will use your default Internet Explorer on your local machine as the embedded browser. If your default browser is Internet Explorer 6, this can cause problems.

If you log in remotely into IBM System Storage Productivity Center, the TotalStorage Productivity Center Element Manager will use Internet Explorer 6 by default as the embedded browser.

There are two work arounds identified:

- Upgrade to Internet Explorer 7. Internet Explorer 7 has improved performance, security, and has this problem fixed.
- Adjust Internet Explorer 6 to not force an HTTP 1.1 connection. Follow these steps:
 1. Open Internet Explorer 6.
 2. Click **Tools > Internet Options**. Click on the **Advanced tab**.
 3. Scroll down to HTTP 1.1 settings, and uncheck the following two items:

- Use HTTP 1.1.
 - Use HTTP 1.1 through proxy connections.
4. Click **Apply**.
 5. Close and then reopen the browser.
 6. You will now be able to access the IBM TotalStorage Productivity Center GUI.

Stop and start the DS Storage Manager Servers

The IBM System Storage DS Storage Manager server and IBM System Storage 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 starting the DS Storage Manager servers on a Windows operating system

You can stop or start 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 (IBM System Storage DS Storage Manager and then click Stop Servers or Start Servers depending on your choice of action. Wait for the servers to start up or completely stop before proceeding. The DS Storage Manager server will stop first and then the DS Network Interface server. The servers start in the opposite order. They can be stopped and started separately, but they must follow this order.
4. Click **Stop**, or **Start** for the action that you want to complete.

Stopping and starting 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. Notice that the DS Storage Manager server stops first and then the DS Network Interface server stops. The servers start in the opposite order.

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 after a 2 hour period of inactivity (no page requests or refreshes) or when you turn off your computer. If your session is inactive for 2 hours, the session will expire and you will be required to sign on again. Any open pages will be closed, and any wizards that have not been submitted will need to be restarted from the beginning. The way you start the DS Storage Manager is different depending on whether you have a Windows or Linux operating system.

Starting the DS Storage Manager on a Windows operating system

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

If you recently restarted the management console, before starting the DS Storage Manager, verify that the DS Storage Manager server and IBM System Storage DS Network server are running. From the Windows Control Panel, select Administrative Tools, then Services.

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

Note: If you have not already done so, you must select the appropriate browser security settings to open DS Storage Manager in a browser. In Internet Explorer, on the Tools menu, click Internet Options. On the Security tab, select Internet and click Custom level. Under Miscellaneous, enable Allow META REFRESH. Under Scripting, enable Active scripting.

1. Click **Start**.
2. Click **Programs**.
3. Click **IBM System Storage 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.
4. Enter the user name and password to access the program. The default user name is **admin**, and the default password is **admin**.

Note: The first time you log on using the default user name and password, a second login screen appears. Change your password on the second login screen. Choose a password that you can readily remember because the password is not available to your administrator nor is it retrievable from the system.

Passwords must meet the following criteria:

- Be six to 16 characters long.
- Must contain five or more letters, and it must begin and end with a letter.
- Must contain one or more numbers.
- Cannot contain the user's user ID.

- Is case-sensitive.
- Four unique new passwords must be issued before an old password can be reused.

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 System Storage DS8000 Signon window is displayed. Enter your user name and password to access the program.

Getting started with the DS Storage Manager

Configuration capabilities and the use of the Copy Services features after the initial installation are dependant on the Operating Environment License storage amounts that you purchased.

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

- The Real-time Manager configuration component (with Copy Services). The real-time DS Storage Manager is available at the management console or through the IBM System Storage Productivity Center.
- The Simulated Manager configuration component

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. You can find additional information about these items in the DS8000 Information Center.

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 **admin**) expires immediately. Use the **Change password** feature to ensure further access to the GUI. Choose a password that you can easily 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 2107/242x is preinstalled with the real-time component of the DS8000 Storage Manager). You can begin using the simulated component immediately after you log into the DS Storage Manager. This component provides the ability to create logical configurations when your storage unit is disconnected from the network. After you create the configuration, you can save it and then apply it to a network-attached DS8000 at a later time.

Attention: The simulated DS Storage Manager is designed to assist you in creating an initial configuration for your system. Do not attempt to upload a configuration file if your system has already been configured.

Real-time component (preinstalled on 2107/242x)

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. The real-time DS Storage Manager is available at the management console or through IBM System Storage Productivity Center.

Activate your license activation codes

These codes must be activated before any logical storage configuration can be applied to your network. The following are examples of licenses that can be activated depending on your purchase:

- Operating environment license for each storage unit that you own. (This license must be activated.)
- Copy Services, multiple codes

Copy Services

You can use the Copy Services selections of the DS8000 Storage Manager if you have purchased the optional Copy Services features.

Activating licensed functions

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

To activate your licensed functions, 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 (Models 921, 922, 931, and 932) or reboot of the affected image (Models 9A2 and 9B2):

- 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

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

Before you can connect 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. Use the diskette if you cannot access the IBM System Storage 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 manually enter them into the system that runs the DS Storage Manager.
- The machine serial number, model, and signature.

To obtain your activation codes, perform the following steps:

1. 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>.
2. 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 or 239x license authorization record when you enter the DS8000 (2107 or 242x) serial number and signature. However, if the 2244/239x license authorization record is not attached to the 2107/242x record, you must assign it to the 2107/242x record in the DSFA application. In this situation, you will need the 2244/239x serial number (which you can find on the License Function Authorization document).

Applying activation codes

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

Notes:

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

activating activities are disruptive and require a machine IML (Models 921, 922, 931, and 932) or reboot of the affected image (Models 9A2 and 9B2):

- Removal of a DS8000 licensed function to deactivate the function. Contact your IBM service representative to perform this operation.
 - 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.
2. Before you begin this task, you must check the error log and resolve any current DS8000 problems. Contact IBM Support for assistance in resolving these problems.
 3. Before you configure, disable or provide paths through any firewalls, because they might interfere with DS8000 communication.

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. Ensure that the Import activation codes page is not open. 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.
2. In the navigation, select **Real-time manager** → **Manage Hardware** → **Storage images**. On the Storage images main page, select the storage image that you want to import the activation code for.
3. From the **Select Action** menu, select **Apply activation codes**. The Apply Activation codes: Real-time 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.

Notes:

- a. For Models 9A2 and 9B2, you need to perform this process for both images, one image at a time.
 - b. The **Authorization level** and **Scope** fields are populated based on the information that is contained within the activation codes.
4. Click **OK** to complete the process.

Importing activation codes

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

Notes:

1. Before you begin this task, you must resolve any current DS8000 problems. Contact IBM Support for assistance in resolving these problems.
2. Before you configure, disable or provide paths through any firewalls, because they might interfere with DS8000 communication.

Perform the following steps to import your activation codes:

1. In the navigation, select **Real-time manager** → **Manage Hardware** → **Storage images**. On the Storage images main page, select the storage image that you want to import the activation code for.
2. From the **Select Action** menu, select **Apply activation codes**. The Apply Activation codes: Real-time page is displayed.
3. Click **Import key file**. The Import page is displayed.
4. In the **Select file to import** field, specify the target file. Use the **Browse** button to navigate to the appropriate directory.
5. After you have specified the code file, click **OK** to complete the process.

Chapter 5. IBM System Storage Management Console

The IBM System Storage 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 also can be connected to your network to enable centralized management of your system using the IBM System Storage DS command-line interface or the storage management software that uses the IBM System Storage DS open API.

The first management console in a storage complex is always internal to the 2107/242x. 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 2107/242x). 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 2107/242x). This console can be cabled together with the primary internal management console of the first 2107/242x.

The management console is the workstation the DS Storage Manager software is installed on. This console is used to configure and manage the DS8000 series system.

Logging onto the management console

The management console must always be left operational. The real-time DS Storage Manager is available at the management console or through the IBM System Storage Productivity Center. 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, if necessary, and turn on the display. 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 the management console session, if you are using a local console, close the lid and slide the keyboard and display back into the rack. You can also choose to logout. Logging out the user returns you to the login prompt without shutting down. This is an appropriate action when you leave the management console unattended.

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 entering 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 using the Hardware Management Console (HMC) located inside the DS8000 or the IBM System Storage Productivity Center (SSPC). The locally installed Web browser on the HMC allows users access to the DS8000 Storage Manager interface without using SSPC. When you use the HMC, 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 HMC.

1. Log on to the WebSM interface at the management console.
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 System Storage DS Storage Manager" to access the DS Storage Manager.

Dual Hardware Management Consoles

An external management console is available as an optional feature and can be used as a redundant management console for environments with high-availability requirements.

The loss of a single hardware management console configuration will not prevent the DS8000 from working. The DS8000 will not be able to call home, change its configuration, and start or stop remote Copy Services through the GUI, CLI, or API. When either an external or internal management console is paired with a single internal management console controlling a single DS8000, this is referred to as a 2X1 configuration. This dual configuration is appropriate for customers with high-availability requirements (particularly with remote Copy Services requirements).

Internet Protocol Conventions for IPv4 and IPv6

The DS8000 supports the use of Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6). Customer network ports can be configured using the hardware management console (HMC) to use either or both standards. IPv6 provides a higher level of security than IPv4.

An IPv4 address is 32 bits. An IP Address is shown as 4 decimal numbers representing 4 bytes: d.d.d.d where d = decimal number (0 - 255). High order bits are the network identifier and lower order bits are the host identifier. The number of bits in a network identifier is defined by a subnet mask which looks like an IP address with all 1s in high order bits and all 0s in low order bits (for example, 255.255.255.0). An IP address with the host identifier set to all 1s is a broadcast address for all hosts on the network. An IP address can also include a port number. The port number follows the IP address and is separated by a colon (for example, 250.250.250.1:8451).

An IPv6 address is 128 bits. The preferred IP address format has 8 hexadecimal numbers representing 16 bytes (for example, x:x:x:x:x:x (x = X'0' - X'FFFF')). An IPv6 address can contain an imbedded IPv4 address. The format has 6 hexadecimal and 4 decimal numbers representing 16 bytes (for example, x:x:x:x:x:d.d.d.d (x = '0'x - X'FFFF' and d = 0 - 255)). A compressed form allows one string of 0s to be replaced by '::' (for example, FF01:0:0:0:0:0:1A2 can be represented by FF01::1A2). High order bits are for the subnet prefix. Lower order bits are for the interface identifier.

An IPv6 address can have a number of bits in a subnet prefix that are defined by a decimal prefix following the IP address (for example, FF01::101/96 has a 96-bit subnet prefix and a 32-bit interface identifier). Most IPv6 address are required to have a 64-bit prefix length. Specific high order bit combinations in the subnet prefix are defined for specific purposes. An IPv6 IP address can also specify an IP port number (for example, [x:x:x:x:x:x]:p, where p = decimal number).

Multiple IPV6 addresses can be assigned to an interface:

Link-local address

Assigned on a given subnet, only one is one required. The first 10 bits of Link local addresses are limited to B'1111 1110 10' (for example, FE80::1)

Global uni-cast address

Used for cross subnet identification. Any prefix other than 1s is reserved for a Link-local address.

Multi-cast, unspecified, or loop back 0 - N any-cast addresses

Used to deliver a message to all interfaces with that multi-cast address. The first 8 bits of a multi-cast addresses is limited to B'1111 1111' (for example, FF01::1).

Transient addresses

Addresses that can also be dynamically assigned for special purposes.

Special addresses**Unspecified address**

Indicates the absence of an address.

Loop back address

The loop back address is ::1 (the host sends a message to itself).

Imbedded IPv4 addresses

Addresses are limited to ::FFFF:d.d.d.d

No broadcast addresses

Addresses that do not have the host identifier set to all 1s.

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.

User Groups

User groups (or roles) are a level of access that is assigned by the administrator, which allows users to perform certain functions. User groups are created using the DS Storage Manager or the CLI.

When a user account is created, the administrator must specify an initial password for the account. This initial password expires immediately which means that the account users must change the password before they are allowed to perform any other actions. This is also true for all account roles, including Administrators.

The user must be assigned to at least one group or role. Users can be assigned to multiple groups or combinations of groups. Groups with the label No Access (only) cannot be selected in combination with another group.

Administrators can make the following user group assignments (Table 3 on page 92 provides specific capabilities for each user group):

Administrator (only)

Must be the only assigned group. This user group has the highest level of authority. It allows a user to add or remove user accounts. This group has access to all service functions and DS8000 resources.

Physical operator (only)

Must be the only assigned group. This user group allows access to resources that are related to physical configuration, including storage complex, storage unit, storage image, management console, arrays, ranks, and extent pools. The physical operator group does not have access to security functions.

Logical operator

Can be assigned in combination with the Copy Services operator group, but not in combination with any other group. This group has access to resources that relate to logical volumes, hosts, host ports, logical subsystems, and volume groups, excluding security functions.

Copy Services Operator

Can be assigned in combination with the Logical operator group, but not in combination with any other group. This group has access to all Copy Services service functions and resources, excluding security functions.

Monitor (only)

Must be the only assigned group. This group has access to all read-only, nonsecurity service functions and all DS8000 resources.

No Access (only)

The default selection. Must be the only assigned group. This group has no access to any service functions or DS8000 resources. This is the user group that is assigned to a user account that is not associated with any other user group.

Table 3. User Group capabilities

Capability	Administrator	Physical Operator	Logical Operator	Copy Services Operator	Monitor	No Access
User account management	X					
Access audit log	X					
Update storage complex	X	X				
Power on/off storage image	X	X				
Update storage unit	X	X				
Update storage image	X	X				
Warmstart storage image	X	X				
Manage arrays, ranks, extent pools	X	X				
I/O port configuration	X	X				
Configuration recovery services (unfence volumes, discard pinned tracks, repair ranks,...)	X	X				
Host configuration	X	X	X			
Logical subsystem configuration	X	X	X			
Volume configuration	X	X	X			
Add or remove volume group	X	X	X			
Assign or unassign volume group to host connection	X	X	X			
Add or remove volumes to volume group	X	X	X			

Table 3. User Group capabilities (continued)

Capability	Administrator	Physical Operator	Logical Operator	Copy Services Operator	Monitor	No Access
Manage Copy Services (FlashCopy, remote mirror and copy, Global Mirror)	X	X		X		
Set Copy Services timeout values	X	X		X		
Update user account password	X	X	X	X	X	X
Query FRUs and enclosures	X	X	X	X	X	
Query configuration	X	X	X	X	X	
Query Copy Services	X	X	X	X	X	
FRU management	X					
Problem management	X					
Validate communication paths	X					
Activate code load	X					
Create a new PE package	X					
Manage storage unit IP addresses	X					

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. 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:
 - Be six to 16 characters long.
 - Must contain five or more letters, and it must begin and end with a letter.
 - Must contain one or more numbers.
 - Cannot contain the user's user ID.
 - Is case-sensitive.

- 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.

The password and its use must meet the following criteria:

- Be six to 16 characters long.
 - Must contain five or more letters, and it must begin and end with a letter.
 - Must contain one or more numbers.
 - Cannot contain the user's user ID.
 - Is case-sensitive.
 - Four unique new passwords must be issued before an old password can be reused.
1. Under **Monitor system**, select **User administration**. In User administration — Main page, select **Modify** in the **Select Action** drop-down box. 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. 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. 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 Administrator 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 Administrator 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 Administrator 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. 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. Beyond the set number of allowable attempts with the wrong password, the account is locked. If the administrative account is locked, the administrator must use the security recovery utility tool to reset the password to the default (administrative). 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, and then create a new password. See DS CLI documentation for more information.

Notes:

1. This security recovery utility tool only unlocks the administrative account on the DS Storage Manager on which the tool is run.
 2. This task only explains how to use the security recovery utility tool to unlock the administrative account. The topic "Unlocking a user account" describes how to unlock a non-administrative user account.
 3. The security recovery utility tool is a script that is installed in a file directory. You run the script from the directory.
1. Press the **Enter** key. The script runs and the administrative account is unlocked. The password is reset to the default (admin).
 2. Open the Web-based System Manager tool and enter the IP address of the HMC that you want to unlock the password for. Press **Enter**.
 3. Enter the HMC (hardware management console) password.
 4. Navigate to the **Service Applications** folder and select **Service Focal Point**.

5. Select **Service Utilities** from the Tasks list, and then highlight the storage facility.
6. Navigate to **Selected** → **Get HMCs**, and then highlight the HMC.
7. Navigate to **Selected** → **Start/Stop ESSNI**, and then select **Reset** under the **Security Recovery Options**.
8. Click **Apply**.

Chapter 7. Storage system logical partition overview

A storage system logical partition (LPAR) 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 9A2/9B2. It allows you to have two virtual storage systems in one physical machine. Each storage system logical partition can run different levels of licensed machine code. Typically, the different machine code levels of a common release are compatible. However, certain levels of licensed machine code might not be compatible if the later licensed machine code level contains substantial enhancements. If your levels of licensed machine code are not compatible, you can contact IBM support to request instructions for installing the two different licensed machine code levels. In addition, if each logical partition runs a different level of machine code, you can contact IBM support to request a code level upgrade on either partition. Your requests will be evaluated and approved based on the feasibility with your particular configuration.

The default resource allocation for processors, memory, and I/O slots in the two storage system logical partitions on the DS8300 is divided into a fixed ratio of 50/50 percent. If variable partitions are required, the resource allocation for processors and memory can be changed to 25/75 percent or 75/25 percent. See “Variable logical partition implementation” on page 98 for more information.

Note: Each storage system logical partition can have 255 logical subsystems (LSSs) with LSS IDs 0–254.

DS8300 LPAR implementation

The DS8300 Model 9A2/9B2 provides a split between the resources.

Each storage image uses the machine type/model number/serial number of the DS8300 Model 9A2/9B2 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.

Figure 8 on page 98 illustrates the default DS8300 LPAR resource allocation.

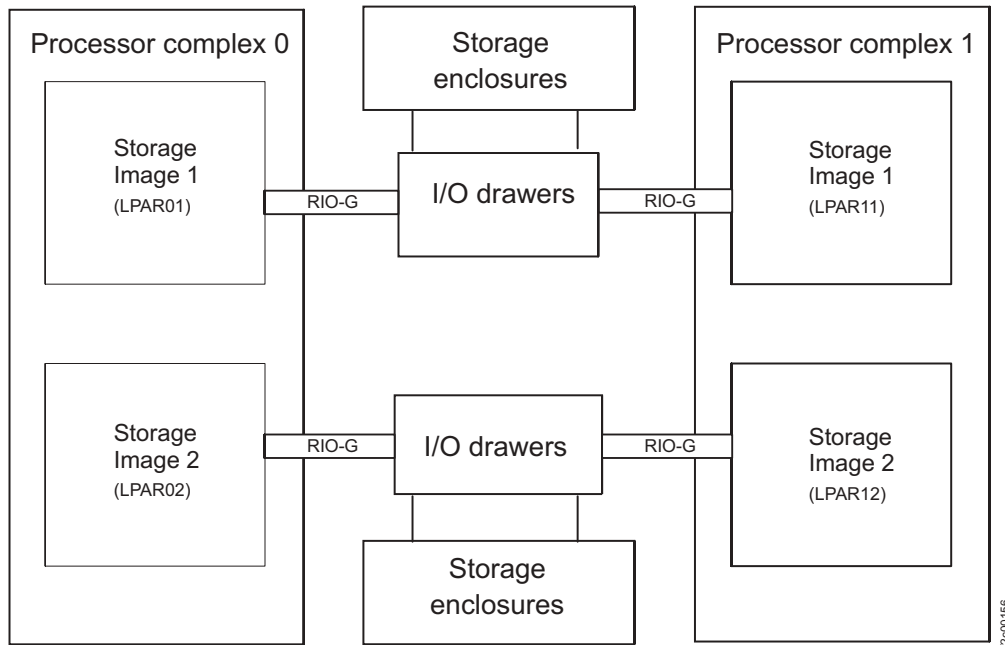


Figure 8. DS8300 default LPAR resource allocation

Each storage image has access to the following components:

- 50 percent of the processors is the default. With variable LPAR implementation on DS8300 Model 9B2, the split can be 25/75 percent or 75/25 percent.
- 50 percent of the processor memory is the default. With variable LPAR implementation on DS8300 Model 9B2, the split can be 25/75 percent or 75/25 percent.
- 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)

Variable logical partition implementation

The DS8300 Model 9B2 can be reconfigured into variable logical partitions (LPARs) that allows for splitting hardware resources where one image has more processor power and cache memory than the other. Variable LPAR implementation allows for increased flexibility, storage image independence, and efficient work load balancing.

The factory default configuration is a 50/50 percent split. Variable LPAR implementations can be re-configured into a 75/25 percent or 25/75 percent split. The reconfigured split applies to the following hardware resources:

- Processors
- NVS+Cache memory
 - Available memory of NVS and cache is split between images
 - Processor memory (equal to installed processor memory features)
 - Image NVS memory (fixed size based on processor memory and the split)
 - Image cache memory (minimum cache memory is 64 GB)

The following hardware resources remain at a 50/50 percent split after the image is re-configured to a 75/25 percent or 25/75 percent split:

- Host adapter slots.

- Device adapter slots.
- Storage enclosure slots.
- RIO-G bus bandwidth

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 System p host between LPARs. However, the DS8300 comes preconfigured with all required LPAR resources assigned to either storage image.

Figure 9 shows the split of all available resources between the two storage images. Each storage image has 50% of all available resources.

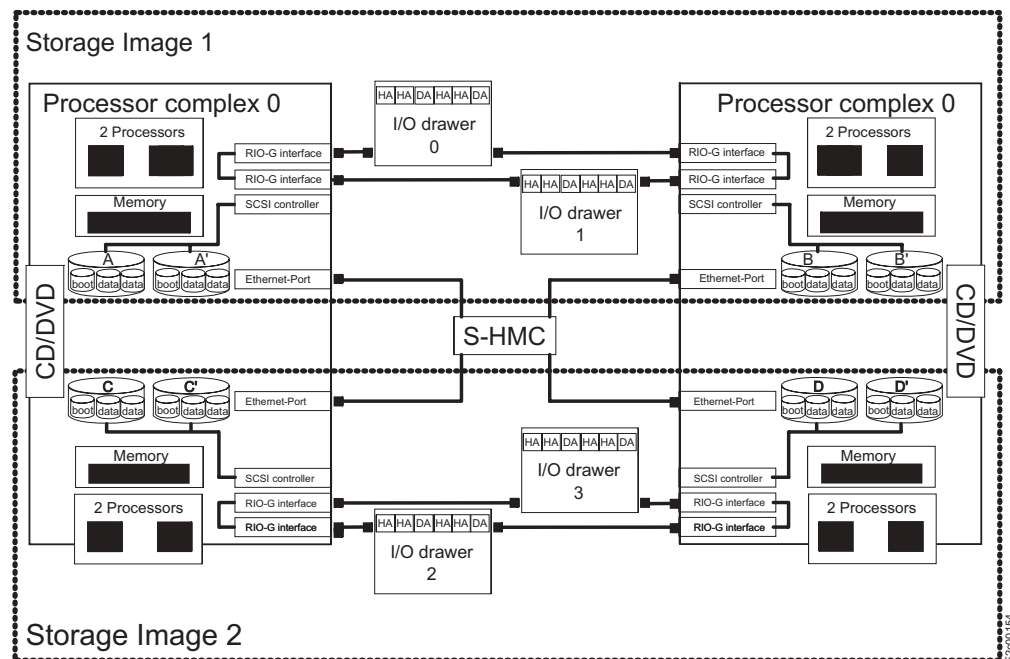


Figure 9. Storage image resource allocation in the processor complexes of the DS8300

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 9A2/9B2, there are four drives in one physical processor complex.

Processor and memory allocation

In the DS8300 Model 9A2/9B2 each processor complex has four processors and up to 128 GB memory capacity. There is also a default 50/50 percent split for processor and memory allocation. The split can be re-configured with variable LPAR implementation. See “Variable logical partition implementation” on page 98 for more information. 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 exists:

- 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 System z or S/390 environments
- Advanced Copy Services functions

Parallel Access Volumes

This feature allows multiple System z or S/390 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 System Storage Resiliency Family. These Copy Services functions offer a variety of data duplication, data migration, and disaster recovery solutions that are supported on open systems and System z or S/390 hosts.

The following features are key components:

- Point-in-time copy, also known as FlashCopy
- Remote Mirror and Copy, previously known as Peer-to-Peer Remote Mirror (PPRC)
 - Metro Mirror, previously known as Synchronous PPRC
 - Global Mirror, previously known as Asynchronous PPRC
 - Global Copy, previously known as PPRC Extended Distance
 - Metro/Global Mirror, a combination of Metro Mirror and Global Mirror functions
- 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 System z or S/390 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 System z and S/390 hosts

This section provides the following guidelines for using parallel access volumes (PAVs) in a System z or S/390 environment:

- 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 when you use a System z or S/390 host.

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 PAV associates the PAV-base address and its PAV aliases in a predefined and fixed method. That is, the PAV-aliases of a PAV-base address remain unchanged. Dynamic PAV, on the other hand, dynamically associates the PAV-base address and its PAV aliases. The device number types (PAV-alias or PAV-base) must match the unit address types as defined in the storage unit hardware.

You can further enhance PAV by adding the IBM HyperPAV feature. IBM HyperPAV associates the volumes with either an alias address or a specified base logical volume number. When a host system requests IBM HyperPAV processing and the processing is enabled, aliases on the logical subsystem are placed in an IBM HyperPAV alias access state on all logical paths with a given path group ID. IBM HyperPAV is only supported on FICON channel paths.

PAV 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 System z or S/390 hosts. 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 System z or S/390 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 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 a System z or S/390 host. It is not used as input to the I/O definition for the operating system for the System z or S/390 host. 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 operating system for the System z or S/390 host requires that the input/output definition file (IODF) that the hardware configuration definition (HCD) creates, specifies base and alias devices. The operating system dynamically discovers 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

Notes:

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 2107/242x for any DS8000 device type. In HCD you can use 3990 as a control-unit type for 3380 and 3390 devices. However, you must use 2107/242x as the control-unit type for 3380B, 3380A, 3390B, and 3390A devices.

Reconfiguring an LCU with larger capacity 3390 volumes

Complete this task 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 LCU

Complete this task to increase the number of 3390s that are formatted within an LCU.

To increase the number of 3390s that are formatted within an LCU, 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.

IBM DS8000 Copy Services functions

The DS8000 supports Copy Services functions that contribute to the protection of data.

These products and features are available on the IBM System Storage DS8000 and IBM System Storage Enterprise Storage Server Models 750 and 800. There are a number of advanced Copy Services functions that are part of the IBM System Storage Resiliency family and are supported across various server environments such as System z, System p, System i, and servers from Sun and Hewlett-Packard.

Note: It is assumed that you have obtained the required information to activate the Copy Services licenses using the IBM Disk Storage Feature Activation (DSFA) Web site at <http://www.ibm.com/storage/dsfa/>. After you obtain the activation keys, it is also assumed that you have entered them in the DS Storage Manager interface.

The following IBM System Storage Resiliency family functions and features are described within this section of the book:

- **FlashCopy** - This Copy Services function enables you to make point-in-time, full volume copies of data, with the copies immediately available for read or write access. It allows you to perform the following types of operations:
 - Reverse FlashCopy relationships
 - FlashCopy to Metro Mirror source volumes
 - Multiple relationship FlashCopy
 - Persistent FlashCopy
 - Resynchronization of FlashCopy volume pairs
- **Remote mirror and copy** - These functions constantly update 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 name is used generically for all methods of synchronous and asynchronous processing. These functions allow you to perform the following operations:
 - Metro Mirror
 - Global Copy
 - Global Mirror
 - Metro/Global Mirror
- **Path connections** - Before you can create source and target remote mirror and 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.
- **zSeries Copy Services** - 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. zSeries Copy Services allows you to perform zSeries Global Mirror and zSeries Metro/Global Mirror operations.
- In addition to the previous functions, the following operations are also associated with Copy Services:
 - Consistency groups

- Failover and failback
- Copy Services functions across a 2105 and 2107
- Copy Services functions with i5/OS
- In addition, the following example Copy Services tasks are included at the end of the chapter:
 - Creating remote FlashCopy transactions (known as inband FlashCopy on the ESS 2105)
 - Adding a Copy Services 2105 domain to a storage complex
 - Creating a Metro Mirror volume pair between a 2107 and a 2105

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.

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, where 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 the 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 want 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 10 on page 110 illustrates how a reverse restore operation works:

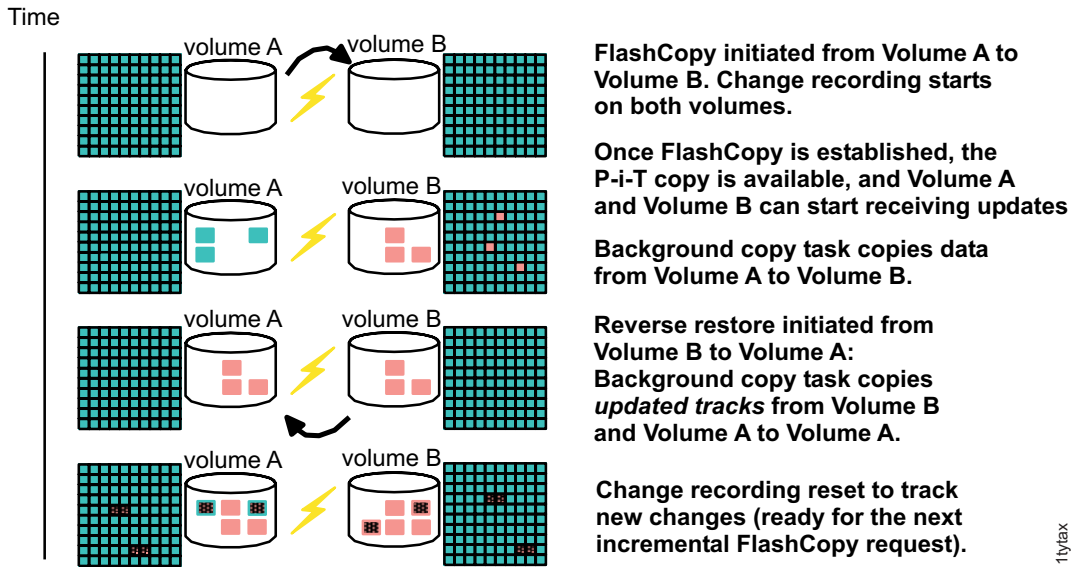


Figure 10. Refreshing target volume — reverse restore

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 the Metro Mirror function.

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) operation of the FlashCopy relationship. The refresh operation 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 the 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 must carefully decide if you can tolerate not mirroring FlashCopy target volumes temporarily at a remote site.

Multiple relationship FlashCopy

The multiple relationship FlashCopy function 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 it is explicitly withdrawn.

Persistent FlashCopy overrides normal FlashCopy processing and allows a FlashCopy relationship to stay active until it is explicitly withdrawn. Normally, 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 identifies the target volume for the most recently completed FlashCopy.

Consider the following situation:

- A company wants to maintain more than one point-in-time version of a set of its production volumes.
- The company creates a FlashCopy version 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 version.

The following represents the procedure that FlashCopy operations follow in this scenario:

1. FlashCopy processing from Volume A to Volume B
2. FlashCopy from processing Volume A to Volume C
3. FlashCopy from processing 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.

Resynchronization of FlashCopy volume pairs

FlashCopy volumes resynchronization (resync) 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 resynchronized. 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 resynchronized, not just the one pair that you selected.

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.

Metro Mirror

Metro Mirror functions offer a synchronous long-distance copy option 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 site 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.

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 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, are copied from the source volume to the target volume. The storage unit updates the target tracks with the current information for each track, regardless of the number of updates between the time that each track 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 processing provides a long-distance remote copy solution across two sites for open systems or System z (or both) data using asynchronous technology. This processing is accomplished using the DS Storage Manager or the DS command-line interface (CLI).

Global Mirror processing is most often associated with disaster recovery or preparing for disaster recovery. However, it can also be used for everyday processing and data migration.

The Global Mirror function is designed to mirror data between volume pairs of a storage unit over greater distances without affecting overall performance. It is also designed to provide application consistent data at a recover (or remote) site in case of a disaster at the local site. By creating a set of remote volumes every few seconds, this function addresses the consistency problem that can be created when large databases and volumes span multiple storage units. With Global Mirror, the data at the remote site is maintained to be a point-in-time consistent copy of the data at the local site.

Global Mirror is based on existing Copy Services functions: Global Copy and FlashCopy. 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.

Your reasons for using Global Mirror processing might include the following:

- Support for virtually unlimited distances between the local and remote sites, with the distance typically limited only by the capabilities of your 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 your local site.
- Data currency, where your remote site might lag behind your local site by 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.
- Session support whereby data consistency at the remote site is internally managed across up to eight storage units 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.

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.

Summary: How Global Mirror works

The automatic cycle in an active Global Mirror session works as follows to maintain data at a remote site to be a point-in-time consistent copy of data at the local site.

1. Consistency groups of volumes are created at the local site.
2. Increments of consistent data are sent to the remote site.
3. FlashCopy operations are performed at the remote site.
4. Global Copy operations are resumed between the local and remote site to copy out of sync tracks.
5. The steps are repeated according to the defined time intervals.

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 System Storage 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 System Storage 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/Global Mirror

The Metro/Global Mirror function combines the capabilities of Metro Mirror and Global Mirror functions for greater protection against planned and unplanned outages.

Metro/Global Mirror is a three-site, high availability disaster recovery solution, which uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site. DS8000 storage units support the Metro/Global Mirror function on open systems and System z or S/390 hosts. You can set up and manage your Metro/Global Mirror configurations using DS CLI and TSO commands.

In a Metro/Global Mirror configuration, a Metro Mirror volume pair is established between two nearby sites (local and intermediate) to protect from local site disasters. The Global Mirror volumes can be located thousands of miles away and can be updated if the original local site has suffered a disaster but has performed failover operations to the intermediate site. In the case of a local-site-only disaster, Metro/Global Mirror can provide zero-data-loss recovery at the remote site as well as at the intermediate site.

In some customer environments, it is necessary to mirror data from a local to a remote site within the distance that is supported for synchronous mirroring. This is especially true when synchronous I/O is required for high or near continuous availability and when a zero-data-loss configuration is required. However, in some cases, it is desirable to have more than a short distance synchronous mirroring solution. Sometimes the following mirroring solutions are required:

- A nearby two-site synchronous copy that can protect from local disasters.
- A longer distance asynchronous copy, at a third site, that can protect from larger scale regional disasters. The third site provides an extra layer of data protection.

The Metro/Global Mirror function provides this combination of synchronous and asynchronous mirroring. Metro/Global Mirror is an extension of Global Mirror, which is based on existing Global Copy (formerly known as PPRC XD) and FlashCopy functions. Global Mirror running at the intermediate site, using a master storage unit and optional subordinate storage units, internally manages data consistency, removing the need for external software to form consistency groups at the remote site.

Figure 11 shows the three sites that are used in a Metro/Global Mirror configuration. The configuration uses a minimum of three storage units, one each at the local, intermediate, and remote sites. A minimum of four volumes (volume A, volume B, volume C, and volume D) are used in this configuration. (For ease of description, the Metro/Global Mirror configuration is described using these volumes. Some environments can contain hundred or thousands of volumes.)

Data from the A volumes at the local site is synchronously replicated to the B volumes at the intermediate site using Metro Mirror. Data from the B volumes at the intermediate site is asynchronously replicated to the C volumes at the remote site using Global Copy. FlashCopy relationships are created with the C volumes at the remote site as the FlashCopy source volumes and the D volumes at the remote site as the FlashCopy target volumes, maintaining the consistent disaster recovery volumes using Global Mirror.

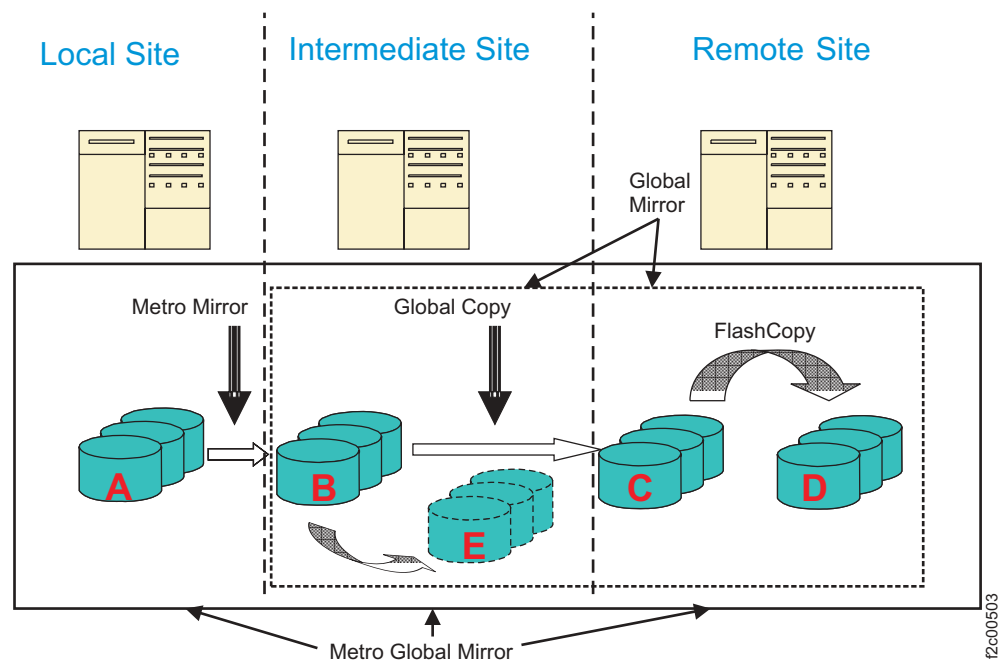


Figure 11. Metro/Global Mirror configuration

As an additional layer of disaster protection if Global Mirror processing were to fail at the remote site, you can use the storage at your intermediate site for a target copy. Setting up Global Mirror between the remote and intermediate sites requires an additional set of FlashCopy volumes at the intermediate site. Then, you can perform failover and restore operations at the remote site using these volumes at the intermediate site (acting as a remote site) to create Global Mirror consistency groups. These volumes, which are referred to as E volumes, would be used as FlashCopy targets for a Global Mirror consistency group.

For Global Mirror processing, one storage unit at the intermediate site is designated as the master storage unit. The master storage unit sends commands over Fibre Channel Protocol (FCP) links and coordinates the consistency group formation process. These links are required for the Global Mirror master storage unit to coordinate the consistency group formation process with the storage units and to communicate the FlashCopy commands to the remote site. All status is relayed back to the master storage unit.

How Metro/Global Mirror works

In order to understand how Metro/Global Mirror works, it is important to understand how the volumes are configured at each site.

Figure 12 shows a Metro/Global Mirror environment that includes the A, B, C, and D volume configuration at the three sites.

There are three volume configurations in a Metro/Global Mirror environment:

- Volumes at the local (or primary) site are Metro Mirror primary (or source) volumes and are referred to as the A volumes.
- Volumes at the intermediate site are the Metro Mirror secondary (or target) volumes and Global Mirror primary volumes at the same time. These volumes are referred to as the B volumes. (If additional volumes are used at the intermediate site as FlashCopy target volumes, they are referred to as E volumes.)
- One set of volumes at the remote site are the Global Mirror secondary volumes and FlashCopy source volumes at the same time. These volumes are referred to as the C volumes. The FlashCopy target volumes are referred to as the D volumes.

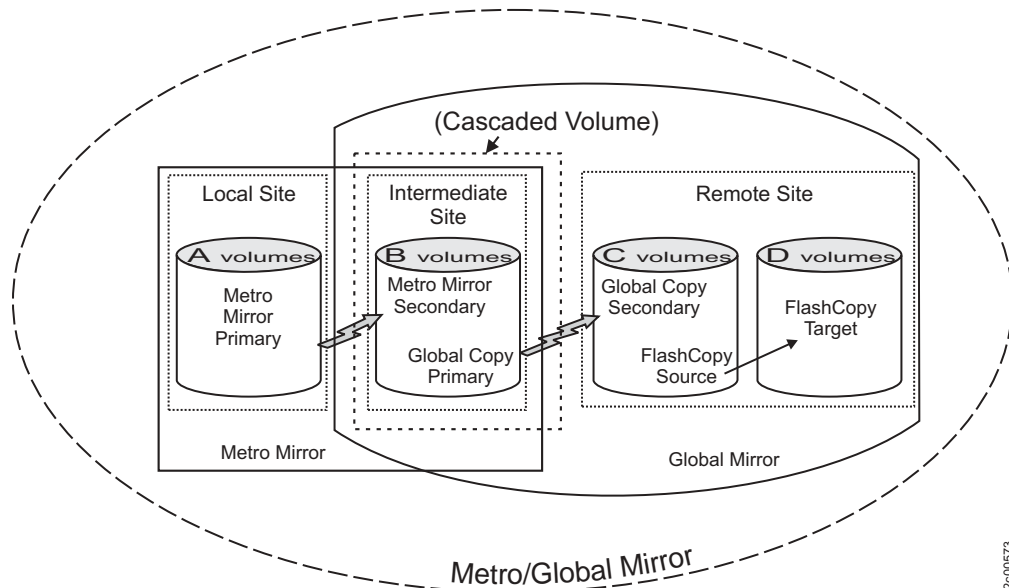


Figure 12. Metro/Global Mirror volume configuration

The Global Mirror function allows you to combine a Global Copy volume pair with FlashCopy such that the Global Copy secondary also serves as the FlashCopy source in a Global Mirror session. FlashCopy plays an integral part of this configuration because the set of volumes that are FlashCopy target volumes contain consistent data. This set of volumes is called a consistency group.

A Metro/Global Mirror environment is configured using a minimum of three storage units, one each at the local, intermediate, and remote sites. Data from the A volumes at the local site is synchronously replicated to B volumes at the intermediate site using Metro Mirror. Data from the B volumes at the intermediate site is asynchronously replicated to the C volumes at the remote site using Global Copy. FlashCopy relationships that are created with the C volumes at the remote site become the FlashCopy source volumes and the D volumes at the remote site become the FlashCopy target volumes.

Note: For ease of understanding, we describe the Metro/Global Mirror configuration in terms of the local, intermediate, and remote sites. You can manage the volumes that are associated with the storage units at these sites from one central location.

For Metro/Global Mirror, one storage unit at the intermediate site is designated as the master. Optional additional storage units at the intermediate site, if configured, are designated as subordinates. The master sends commands to its subordinate storage units over fibre channel protocol (FCP) links and coordinates the consistency group formation process over all the subordinate storage units. The FCP links also enable the Global Mirror master storage unit to communicate the FlashCopy commands to the remote site. All status is then relayed back to the master storage unit. Figure 13 provides a high-level view of the three-site Metro/Global Mirror configuration.

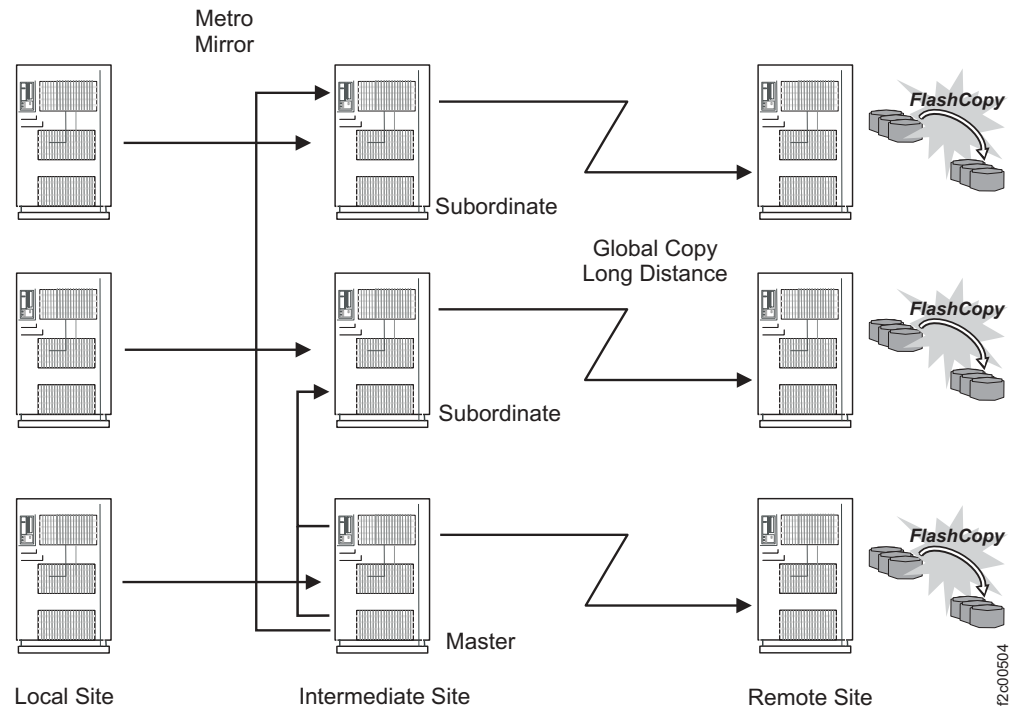


Figure 13. Metro/Global Mirror configuration with master and subordinate storage units

In a Metro/Global Mirror configuration, if you lose access to the storage unit at the local site, the Global Copy relationship is intact and Global Mirror is still operational. Through the use of FlashCopy operations and consistency groups, consistent data is maintained at the remote site. However, you are left without the local site until all recovery operations have been performed to recover the local site.

A three-site Metro/Global Mirror configuration provides the following recovery options at the alternate sites if a failure occurs:

- If an outage occurs at the local site, recovery operations can begin at the intermediate site. Global Mirror continues to mirror updates between the intermediate and remote sites, maintaining the recovery capability at the remote site.
- If an outage occurs at the local site, recovery operations can begin at the remote site and preparations can be made to resynchronize the local site when it recovers from its disaster. Once in recovery mode at the remote site, another Global Mirror session can be set up and put into operation using the former intermediate as its new remote site. This new Global Mirror session provides additional disaster recovery while operating at the remote site.
- If an outage occurs at the intermediate site, data at the local storage unit is not affected. Applications continue to run normally.
- If an outage occurs at the remote site, data at the local and intermediate sites is not affected. Applications continue to run normally. The intermediate storage unit maintains a consistent up-to-date copy.
- If both the local and intermediate sites are lost, the scenario is similar to a two-site scenario when access to the local storage unit is lost. Recovery must be achieved using the last consistent point-in-time copy at the remote site.

Consistency group formation in a Metro/Global Mirror environment

Consistency groups in a Metro/Global Mirror environment create a consistent point-in-time copy of multiple volumes, with negligible host impact.

If there is a disaster at the local site, data consistency must be maintained across multiple volumes. This requires ensuring that consistency groups form successfully at the remote site and are valid to use. (A consistency group keeps data consistent across multiple volumes in multiple storage units.)

There are two methods of maintaining consistent data in a Metro/Global Mirror configuration:

Metro Mirror

The freeze function that is part of Metro Mirror operations provides data consistency for a Metro Mirror environment. The freeze function can be used by external automation software to ensure data consistency to multiple Metro Mirror volume pairs in a specified LSS. External automation software can detect messages that are generated by Metro Mirror processing errors. You can temporarily suspend write operations to the volume pairs that are associated with the specified LSS. This ensures a point of consistency at the remote site, across all necessary volumes in the specified LSSs. Afterwards, you can resume host I/O on the volumes pairs.

Global Mirror

The integrated consistency capabilities of Global Mirror, which uses a master storage unit (along with optional subordinate storage units) internally manage data consistency using consistency groups in a Global Mirror environment.

Note: The consistency of the Global Mirror replication depends on the consistency of the Metro Mirror target volumes. Otherwise, even though the Global Mirror target volumes will be a valid copy of the Global Mirror source volumes during the mirroring process, they might not necessarily be a valid copy of the production volumes.

The process of forming consistency groups using Global Mirror involves the following steps:

1. **The consistency group is created.**
2. **Data in the consistency group is sent to the remote site.**
3. **Data in the consistency group is saved as a group.**

In a Metro/Global Mirror configuration, Metro Mirror and Global Copy volume pairs, combined with consistency groups using Global Mirror and FlashCopy, help maintain consistent and current (usually within seconds) data at the remote site to provide a true backup at the remote site. Metro/Global Mirror introduces Global Mirror and FlashCopy change recording and out-of-sync bitmap features that allow faster subsequent resynchronization of the volumes at the local site.

During a disaster while the intermediate site is in production mode, the change recording feature continues to be used to copy changed data from the B volumes at the intermediate site to the C volumes at the remote site as consecutive sets of consistency groups. When the original local site is ready to resume production mode, the direction of the Metro Mirror configuration must be reversed. The out-of-sync bitmaps are used to resynchronize the volumes at the local site.

During normal operations, consistency groups are continually being created from the volumes at the intermediate site. While consistency groups are being formed, host write I/Os are momentarily delayed while a consistent data point is attained (because Metro Mirror from the local to intermediate site is synchronous). Bitmaps that track updates for the consistency group are created and data is transferred, or *drained*, to the volumes at the remote site using Global Copy. Application writes for the next consistency group are recorded in the change recording bitmap.

The master storage unit coordinates the consistency group formation for all volumes that belong to a Global Mirror session over FCP links that are based on user-specified parameters. It also coordinates I/O with optional subordinates to all primary volumes to form consistency groups. The frequency that consistency groups are formed is based on the number of seconds that is specified for the Consistency Group Interval option and the amount of time that it takes to form consistency groups. Other Global Mirror options are described in “Consistency group tunable values” on page 124.

Note: Be aware that an insufficient bandwidth (caused by a link failure or insufficient link configuration) can also increase the amount of time that it takes to drain a consistency group.

In a Metro/Global Mirror configuration, if a Metro Mirror volume pair between the local and intermediate sites is either duplex pending or becomes suspended, the Global Mirror session will fail consistency group formation. This is to ensure that the consistency group at the remote site is, in fact, truly consistent with the local site.

In addition, there might be cases where the secondary volume in a Metro Mirror relationship at the intermediate site is not aware that its status as a secondary is suspended and the consistency group formation would continue. The Global Mirror session fails the consistency group formation as a way to ensure that the consistency group that is formed at the remote site is consistent with the local site. You can query information about the intermediate volume’s status in the Global Mirror session using the `showgmir -metrics` command.

After the master storage unit has coordinated a consistent point across the source volumes of the Global Mirror pairs (the B volumes in a Metro/Global Mirror configuration), the master storage unit monitors the progress of the transfer of the consistency group to the remote site. When the consistency group transfer has completed, the master storage unit then establishes the FlashCopy relationships at the remote site.

In a Global Mirror environment, some rare error cases can occur where the FlashCopy operations on the remote site are interrupted. The consistency group is protected and can be recovered. If the site where Global Mirror source volumes reside is still operational, then Global Mirror will enter a *fatal* state to preserve the consistency group. In a Metro/Global Mirror environment, this site is the intermediate site.

If the site (also the intermediate site in a Metro/Global Mirror environment) where the Global Mirror source volumes reside is not operational because of a disaster, the consistency groups on the remote site might only be partially formed. In either case, you must verify the consistency group on the remote site and determine whether the changes need to be “rolled forward” (committed) or “rolled backward” (discarded).

Consistency group tunable values

As part of the consistency group formation, the master storage unit must communicate with the subordinate storage units to determine a consistent data point across volumes. Point-in-time copies are written to the tertiary volumes, the FlashCopy target volumes at the remote site. Figure 14 provides an illustration of these values.




Figure 14. Forming consistency groups according to specified intervals

User-specified values for the following parameters control how often consistency groups are formed.

Consistency group interval time

The amount of time that the master storage unit can spend coordinating with the subordinates to form a consistent data point. It is the amount of time that Global Mirror is in Global Copy mode between the formation of

each consistency group. The default is zero seconds, which means that Global Mirror will continuously form consistency groups as fast as the environment will allow.

In most situations, it is recommended that you leave the default value of zero to allow Global Mirror to form consistency groups as fast as possible. Global Mirror automatically moves to Global Copy mode for a period of time if the drain time is exceeded.

Maximum consistency group drain time

The maximum amount of time that Global Mirror spends draining data from the A and B volumes in a consistency group. If this process is not done during this time, a consistency group failure is reported. If the maximum drain time is exceeded, then Global Mirror will transition to Global Copy mode for a period of time in order to catch up in the most efficient manner.

The default maximum drain time value is 30 seconds. Setting a maximum drain time protects production performance by not permitting consistency groups to form at times when it might not be desirable. It is recommended that you not change the default of this setting.

Maximum coordination interval

The maximum time that Global Mirror will allow for the formation of the consistency group before failing this consistency group. Having this time value set ensures that if there is an error recovery event or communications problem, applications will not experience significant impact from consistency group formation failure. The default is 50 ms. It is recommended that you not change the default of this setting.

Automated disaster recovery offerings for Metro/Global Mirror

The Metro/Global Mirror function has a number of supported automated management offerings for disaster recovery solutions.

IBM offers a portfolio of automated solutions that are available to help you implement a disaster recovery plan that best fits your IT environment. These solutions require some sort of automation that can react to a failure in a timely manner. Automated solutions, such as GDPS with Global Mirror, can monitor and manage data resources during planned and unplanned outages.

GDPS is an automated disaster recovery solution for System z environments. In a GDPS environment, a combination of Remote Mirror and Copy and Global Mirror, or HyperSwap™ Manager and Global Mirror, manages the mirroring of data (both disk and tape) between your local and remote sites. Data consistency is provided across all platforms that share the storage units in a GDPS configuration.

When you combine the Global Mirror and Metro Mirror technologies with automation, you can ensure continuous or near-continuous availability during outages, including disasters. With automation, synchronization is achieved through the point in time of the *freeze* by issuing commands in parallel to all logical subsystems (LSSs) that are contained within the Metro Mirror consistency group configuration. This ensures consistent data across all LSSs in the secondary copy of data during a disaster.

With GDPS, you can write scripts or even a *single* GDPS script to perform tasks that allow operations to be run in parallel across systems. The amount of work is spread evenly across the defined tasks in parallel. GDPS can consolidate and check the results and subsequently provide feedback on the status of the operations.

You can also customize tasks for other user-defined actions, such as planned site switches in which the workload is switched from one site to another. For further information about GDPS, see *GDPS/PPRC V3R2 Installation and Customization Guide (or GDPS/PPRC HyperSwap Manager V3R2 Installation and Customization Guide)* and *GDPS Global Mirror V3R2 Installation and Customization Guide*. GDPS is available from IBM Global Services.

Note: The GDPS Installation and Customization Guides are licensed publications available only to customers that are licensed for GDPS. If you are not licensed, you can review the publication at <http://www.redbooks.ibm.com/redpieces/abstracts/sg246374.html?Open> and contact IBM for further details on these offerings.

Incremental resynchronization in a Metro/Global Mirror configuration

The purpose of the incremental resynchronization function is to help you avoid having to do a full volume resynchronization between the local and remote sites if an outage occurs at your intermediate site.

In a three-site Metro/Global Mirror configuration without incremental resynchronization, if there is an outage at the intermediate site, reestablishing a disaster recovery environment between the local and remote sites requires a full volume resynchronization. To maintain data integrity, a full volume copy is performed because there is no guarantee that data that was written to the local Metro Mirror volumes was secured at the remote site when the outage occurred. There is also no guarantee that Global Mirror formed successful consistency groups of consistent data at the remote site before the outage occurred. During this full volume resynchronization period, which can take a long time depending on the amount of data being copied, your data *can* potentially be susceptible to a loss.

To ensure the retention of your data and restart a disaster recovery environment running Global Mirror between the local and remote sites, you can establish the primary Metro Mirror volumes at the local site as the primary volumes to the Global Copy target volumes at the remote site. This implementation, shown in Figure 15, gives you the flexibility to establish asynchronous mirroring between the local and remote site. Consistency groups can continue to form at the remote site, while retaining data consistency over distance between the two sites.

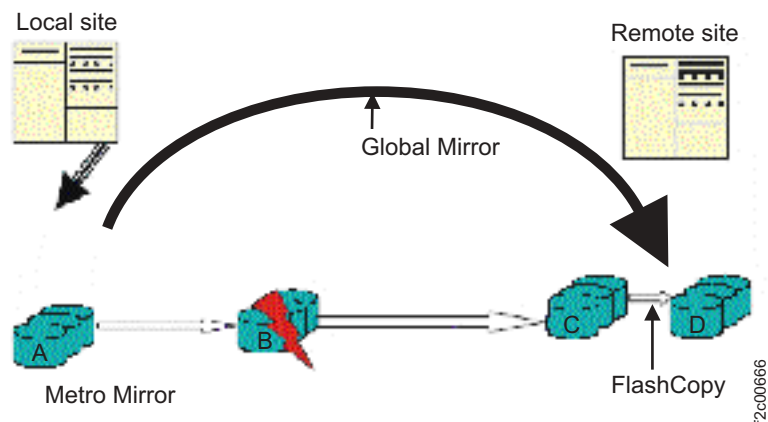


Figure 15. Incremental Resynchronization implementation

The incremental resynchronization function involves a change recording mechanism on the specified primary Metro Mirror volumes in a Metro/Global

Mirror configuration. The volumes are continually updated to maintain a record of tracks that have changed data that has not been secured at the remote site. If the intermediate site fails, the change recording bitmaps and out-of-synchronization bitmaps on the primary Metro Mirror volumes contain the tracks of data that are needed to resynchronize the local site and remote site volumes. The resynchronization process copies only a subset of changed data that has to be copied between the local site and remote site volumes and eventually to the D volumes (the FlashCopy target volumes in Global Mirror).

The incremental resynchronization function uses a change recording mechanism to track in flight changes. Because there can only be one function that actively uses change recording on a volume at a time, the incremental resynchronization function prevents change recording for other purposes, such as incremental FlashCopy processing. If an attempt is made to start incremental resynchronization when another function using change recording is active, that start request fails. If an attempt is made to do an incremental FlashCopy from a remote mirror and copy primary volume that has the incremental resynchronization active, that FlashCopy request fails.

When the intermediate site is operational, you can return to the original cascading Metro/ Global Mirror configuration. You can restore your Metro/Global Mirror configuration without interrupting production.

How Incremental Resynchronization works

If you enable the incremental resynchronization function for Metro/Global Mirror and there is an intermediate site failure, the procedure to restart the recovery environment running Global Mirror from the local site to the remote site is simplified, and a full volume copy is not required.

The following sections describe starting the incremental resynchronization function with initialization in a Metro/Global Mirror configuration and, in some very specific situations, without initialization. Information is also provided about other options that you can use with the incremental resynchronization function.

For best management practices, it is recommended that you manage your Metro/Global Mirror configurations using external software in conjunction with the IBM DS CLI or IBM TotalStorage Replication Manager.

Starting the incremental resynchronization function with initialization

Starting incremental resynchronization with initialization means that the change recording mechanism is started on the specified primary Metro Mirror (A) volumes in a Metro/Global Mirror configuration. This creates an environment on the primary Metro Mirror volumes that enables the microcode to monitor and track data "in flight" in a Metro/Global Mirror configuration. (*In flight* data refers to data that has been written to the primary volumes but has not been secured on the C volumes at the remote site.)

When you set up your Metro/Global Mirror configuration, you must enable the incremental resynchronization function with the initialization option to create and start change recording bitmaps on each of the primary Metro Mirror volumes. When data is written to the primary volumes, an out-of-synchronization change recording bitmap is also created on the primary Metro Mirror volumes, which is later used for resynchronization between the local site and remote site volumes.

Periodically, the incremental resynchronization function performs queries from each of the A volumes at the local site to each of the B volumes at the intermediate site to determine when data has been secured on the C volumes at the remote site.

The queries generate a report that allows you to view the following types of information:

- How often consistency groups are being formed in the Global Mirror environment
- The number of successful consistency groups that have formed since the last query

In addition to the periodic queries, you can issue a query from the DS CLI command prompt using the **showgmir -metrics** command. This command generates a report that displays the status of consistency group formation.

If the Metro Mirror volume pairs are in full duplex status and at least three consistency groups have formed successfully, then data from the local site volumes is secure on the remote site volumes. The incremental resynchronization function uses a mechanism to determine how often and how many tracks must change on the specified primary volumes before it resets (and clears) the bitmaps after in flight data is secured on the remote site volumes. This process continues until you disable the incremental resynchronization function on the primary volumes using the **mkpprc** command or the primary Metro Mirror volumes become secondary volumes.

If you started the incremental resynchronization function with initialization and an outage occurs at the intermediate site, you can restart a recovery configuration with Global Mirror running between the local and remote sites.

Starting the incremental resynchronization function without initialization

Starting incremental resynchronization *without* initialization means that the change recording bitmaps are created but started with all zeros on the primary Metro Mirror volumes. Do *not* specify the incremental resynchronization function without initialization except in some specific recovery scenarios.

Recovering with the incremental resynchronization function

In an existing Metro/Global Mirror configuration, if the intermediate site fails, use the **mkpprc** command with the **-incrementalresync recover** parameter to restart your disaster recovery environment. This initiates the processing of Global Mirror between the local and remote sites.

Before the **mkpprc** command completes its process, the system verifies that the specified volumes have the intermediate volumes in common. This verification ensures that the primary volume specified in the command is the Metro Mirror A volume and the secondary volume specified is the Global Copy C volume of a Metro/Global Mirror configuration.

When the **mkpprc** command processes, only changes on the A volumes that have the change recording mechanism started are copied to the C volumes at the remote site. A full volume copy of these volumes is not required.

Note: When incremental resynchronization is started, there is a short period of time where a full volume copy is required. This occurs because incremental resynchronization has not fully initialized.

Overriding incremental resynchronization function checking

You can override (or bypass) system verification to avoid checking if relationships exist between the volumes in a Metro/Global Mirror configuration. Issue the **mkpprc -incrementalresync override** command to override the system verification. When the command processes, only changes on the Metro Mirror A volumes that have the change recording mechanism started are copied to the C volumes at the remote site. Therefore, you must ensure that the specified volumes have the intermediate volumes in common before processing the command.

You might want to override a relationship verification in specific scenarios. For example, you might have a circumstance where the intermediate site is operational. If this is the case, it is not required to verify whether former relationships exist when you want to return to the original volume pair relationships in the cascading Metro/Global Mirror configuration.

Stopping the incremental resynchronization function

You can stop incremental resynchronization for the specified Metro Mirror A volumes in a cascaded Metro/Global Mirror relationship. Issue either the **mkpprc -incrementalresync disable** command or the **rmpprc** command at the DS CLI command prompt.

Path connections

Before you can create source and target remote mirror and 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.

Copy Services for System z and S/390 hosts

Copy Services functions are also supported in a System z or S/390 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.

Copy Services for System z and S/390 hosts includes z/OS 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 System z data using asynchronous technology.

DS8000 storage complexes support the z/OS Global Mirror function only on System z hosts. The 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 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 synchronized again between the copies.

z/OS Metro/Global Mirror

z/OS Metro/Global Mirror is an optional function that is available with the remote mirror and copy feature for System z hosts and is supported on the DS8000.

The mirroring capability of z/OS Metro/Global Mirror uses 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 solution for even greater protection from unplanned outages.

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 that 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. The Global Copy target volumes at the recovery site are likely to be inconsistent. Consistency of the volume pairs can be ensured 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).

Copy Services functions across a 2105 and 2107/242x

Copy Services functions that are performed using either the DS Storage Manager or DS CLI are designed to interact with both the 2107/242x and the IBM TotalStorage Enterprise Storage Servers (ESS) Models 750 and 800.

Most Copy Services functions that are available on the ESS 2105 are also available in the 2107/242x and in open systems and System z environments.

On the 2107/242x, 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 2107/242x at your local (or primary) site and with the 2105 at your secondary (or recovery) site, where the same performance is not required.

Consider the following guidelines:

- To perform Copy Services functions between machine types 2105 and 2107/242x, 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 2107/242x. 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.
- To manage Copy Services across the 2105 and 2107/242x, you must install licensed internal code version 2.4.2 or later on the ESS Model 800.
- The 2107/242x does not 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 2107/242x and 2105, you must use a FCP link.

Copy Services functions with i5/OS

You can use Copy Services functions with IBM i5/OS to perform point-in-time copies of your data using the FlashCopy function and mirroring your data for disaster recovery using remote mirror and copy functions.

General considerations with i5/OS (prerequisites for Copy Services functions)

Many initial configuration tasks to use i5/OS with the DS8000 are the same across the various operating systems. However, unlike other operating systems, there are some storage requirements that are unique to i5/OS. For example, LUNs on i5/OS are defined as fixed block storage and only specific LUN sizes are supported, such as 8.5 GB, 17.5 GB, and 35.1 GB. i5/OS LUNs make available a 520-byte block to the host. i5/OS refers to volumes as LUNs.

There are two types of LUN protection: protected and unprotected. In reality, both volume protection types have the same DS8000 internal RAID protection. If you define a LUN as protected, it is protected by the hardware (the DS8000) rather than by i5/OS host-based mirroring. If you define a LUN as unprotected, either the i5/OS or another LUN of equal capacity, either internally or externally mirrors it. The unprotected type is typically used for Load Source Unit mirroring. (Each System i or LPAR partition has a special disk called the Load Source Unit. It holds the system's microcode and is used to reboot the system.) Unless you intend to use i5/OS (host based) mirroring, you should define your LUNs as protected.

In addition, i5/OS supports multipathing, which is part of the base operating system. Multipathing provides redundant path solutions for System i but is only used for external storage like the DS8000. Each connection for a multipath disk unit functions independently. Several connections provide availability by allowing disk storage to be used even if a single path fails.

Some i5/OS considerations are similar to other operating systems regarding how LUNs are defined to volume groups. For example, to access volumes for Copy Services functions with i5/OS, a separate volume group for the target volume IDs is required. Because you do not want to write to the target volumes from the same host while Copy Services functions are running, you might want to put the target volumes in an unassigned volume group until the volumes are needed.

Assuming that you have set up and configured your DS8000 for i5/OS, you can perform FlashCopy and remote mirror and copy operations to create copies of System i disk pools within a single DS8000 or to a distant DS8000 using remote mirror and copy functions. You can save this data and back it up to tape.

You can install and run the DS CLI on i5/OS to allow you to perform Copy Services functions that are supported on the DS8000. You can also use the DS Storage Manager for i5/OS to perform Copy Services functions. The size of the target volumes must be equal the size of the corresponding source volumes. Any mismatch causes the task to fail.

System i Copy Services Toolkits

The System i Copy Services Toolkit allows you to use the data mirroring capabilities of Copy Services functions on the DS8000. The Toolkit allows you to take advantage of these functions with minimal downtime of your production system. The Toolkit interfaces with the DS CLI and DS Storage Manager to send Copy Services commands to the DS8000. For more information about the Toolkit, contact your IBM representative.

FlashCopy with the System i Copy Services Toolkit

Once configured and implemented, the Toolkit will manage all of the System i functions that are necessary to use FlashCopy, which can involve creating a point-in-time version of an independent disk pool called an independent auxiliary storage pool (IASP).

Note: In a System i environment, disk pools are included in auxiliary storage pools (ASPs). Data is scattered across all disks in a given ASP. The first ASP on the system is called the System ASP and includes the operating system. Users can optionally have additional ASPs called user ASPs. These are typically used to hold database journals and savefiles, since these objects need to be separated from

the objects they are protecting. The System ASP and all the user ASPs together are called traditional ASPs.

An independent auxiliary storage pool is an independent disk pool. It is a method of configuring multiple disk units into one group. (Disk pool, IASP, and database are sometimes used interchangeably.) An IASP can be varied online or offline together independent of the rest of the storage on a system, including the system ASP, user ASPs, and other independent disk pools. After IASPs are implemented, the FlashCopy operation can be applied to just the IASP, not the entire system.

Remote mirror and copy with the System i Copy Services Toolkit

Like the Toolkit for FlashCopy, once configured and implemented, the remote mirror and copy Toolkit manages all of the System i functions that are necessary to allow you generate a second copy of the IASP on another storage unit.

Backup Recovery and Media Services with FlashCopy

Backup, Recovery and Media Services (BRMS) for System i is software that helps manage your backups and provides you with an orderly way to retrieve lost or damaged data.

Using BRMS, you can manage your most critical and complex backups simply and easily. You can also recover your system fully in the event of a disaster or failure. In addition to these backup and recovery features, BRMS enables you to track all of your backup media from creation to expiration. You no longer have to keep track of which backup items are on which volumes or worry that you will accidentally write over active data.

The BRMS implementation of FlashCopy provides a way to perform a backup on a system that was copied using the FlashCopy function. BRMS history appears like the backup was performed on the production system.

For more information about BRMS, visit the i5/OS Information Center on IBM Publications Web site at <http://publib.boulder.ibm.com/eserver/>.

Adding a Copy Services 2105 domain to a storage complex

Complete this task 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 2107/242x. The mirroring solutions are compatible between the 2105 (ESS Models 750 and 800) and 2107/242x.

For this example, assume that you want to configure a 2105 Copy Services domain using the DS Storage Manager. The 2107/242x must authenticate with the 2105 before you can perform Copy Services functions. Because of the single login process used by the 2107/242x to access the 2105, the user names and passwords

must be the same on the 2107/242x and 2105. This allows the user to access the 2105 server from the DS Storage Manager, without having to authenticate on the 2105.

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**.
 - 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 2107/242x .

Creating a Metro Mirror volume pair between a 2107 /242x and a 2105

Complete this task to create a Metro Mirror volume pair using volumes from a 2107/242x 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 machine types 2107/242x and a 2105, you must have added the **Add 2105 Copy Services Domain** option from the Select 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 2107/242x must be configured using Fibre Channel Protocol (FCP) ports.
- The storage type of the source and target volumes on the 2107/242x 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 2107/242x and a 2107/242x
- A 1750 and a 1750
- A 2107/242x and a 1750
- A 2105 and a 1750
- A 2105 and a 2107/242x

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.

Perform the following steps to create a Metro Mirror pair between a 2107/242x and a 2105. For this task, the source domain is a 2105 Model 800 or 750 and the target is a 2107/242x.

1. In the navigation, select **Real-time Manager** → **Copy Services** → **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 Select 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...**
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 2107/242x 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.

Creating remote FlashCopy transactions

Complete this task to create a remote FlashCopy (inband FlashCopy on the ESS 2105) at a target (remote) site using remote FlashCopy commands.

Remote FlashCopy operations can only be processed using the DS CLI and not the DS Storage Manager. (Part of the Remote FlashCopy operation requires that you create paths first. You can issue those requests using either the DS Storage Manager or the DS CLI.)

To establish a FlashCopy relationship at 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. This eliminates the need for a network connection to the target site solely for the management of FlashCopy relationships.

Limitation: Remote FlashCopy commands establish a FlashCopy relationship at the target (remote) site when a network connection to the target site is lost. The Remote FlashCopy operation is not supported through the DS Storage Manager, because network connections to both the source and target sites are required. If the network connection to the target site is lost, the DS Storage Manager cannot connect to the target site. Whether you use the DS Storage Manager or the DS CLI for Steps 1 and 2, you must perform Step 3 from the DS CLI.

Note: You can perform all steps from the DS CLI.

Perform the following steps to create a remote FlashCopy operation:

1. **Create paths between the source LSS and the target LSS.** For example, IBM.2107-1300861 and IBM.2107-1300871. You need to know which volumes are available for use before you can issue the request to establish the path.
2. **Create Metro Mirror volume pairs from the source LSS to the target LSS.** For example, volume 2200 (IBM.2107-1300861/0001) from LSS22 and volume 2A00 (IBM.2107-1300871/0001) from LSS22.
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 network connection is lost. You can create the FlashCopy relationship from volume B to volume C (both volumes at the target site). However, you cannot use the DS Storage Manager for this step because connections to the target site are lost. You must use the DS CLI for this step (see Processing Remote FlashCopy [inband] transactions in the *IBM System Storage DS8000 Command-Line Interface User's Guide* for additional information.

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
- Linux for Intel
- OS/400
- System i (i5/OS)
- System z (z/OS, z/OS.e, z/VM, VSE/ESA, UNIX, and TPF1)
- TRU64 UNIX
- Windows Server 2000 and Advanced Server with SP4
- Windows Server 2003
- Windows Server 2008
- Windows XP with SP1

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.

Reformatting disk drive modules on open systems

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 a storage unit, the format is a fixed-byte sector of 524-bytes. The data portion of the sector remains at 512-bytes. System i 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 a 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.

System i 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 **mklvcopy** command:

```
# mklvcopy -e m -s y -k data1v 2 hdisk3 hdisk7
.
.
.
# splitlvcopy -y splitlv data1v 1
```

The **mklvcopy** 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 **mklvcopy** command starts mirroring the `infx1v` logical volume.

```
# mklvcopy -e m -s y -k infx1v 2 hdisk4 hdisk10
.
.
.
# rmlvcopy infx1v 1 hdisk4
```

In this example, `hdisk4` is where the data for `infx1v` 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 IBM System i host system to the DS8000

Use the following information as a guide to determine how to migrate data from a System i host system to the DS8000.

When you create the System i 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.

The DS8000 supports 2107/242x device types with model A81-A87 (unprotected) and A01-A07 (protected) and you need a level of i5/OS or 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).

System i 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 System i host. Up to a maximum of 32 LUNS per

fibre-channel adapter are supported by the DS8000. If the multipathing option (only available on i5/OS V5R3 and above) is selected, then a maximum of 16 LUNs per adapter is recommended.

Because the System i 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 System i volumes are RAID 5 and are protected. When you create the System i 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 IBM System z host system to the DS8000 storage unit

You can use various methods to migrate data from a System z 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
Linux on System z environment	IBM System Storage Global Mirror , Remote Mirror and Copy (when available)
S/390 environment	IBM System Storage Global Mirror , Remote Mirror and Copy (when available)
System z9™ environment	IBM System Storage Global Mirror , Remote Mirror and Copy (when available)
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.

Environment or Operating System	Data Migration Method
z/OS operating system	<ul style="list-style-type: none"> • DFSMSdss (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 • Softek Transparent Data Migration Facility (TDMF), available from IBM as a data migration service
z/VM operating system	<p>DFSMS/VM with the following functions:</p> <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.

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 disk drive modules (DDMs). 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 **mklvcopy**, **splitlvcopy** method

This method is ideal for creating a copy of logical volumes. Use the **mklvcopy** command to ensure data integrity by creating a mirror copy of the data and synchronizing it automatically. Do not run the **splitlvcopy** 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 **splitlvcopy** 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 **mklvcopy**, **rmlvcopy** 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 **splitlvcopy** 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 host system

Use the following information as a guide for how to migrate data from a Windows Server 2000, 2003, or 2008

You can migrate data from a Windows host system 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 System Storage 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 System Storage Metro Mirror feature.

3. Convert the pairs to IBM System Storage 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 System Storage 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 System Storage Metro Mirror.

If you use Global Copy, convert the pairs to IBM System Storage 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 System Storage 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 System Storage 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 System Storage Metro Mirror.
3. If you use Global Copy, convert the pairs to IBM System Storage 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 to prepare for a possible disaster by backing up, copying, and mirroring your data both at the local (production) and remote sites.

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. Consistent data to the same point-in-time across all storage units is vital before you can recover your data at a backup (normally your remote) site.

Most users use a combination of remote mirror and copy and point-in-time copy (FlashCopy) features to form a comprehensive enterprise solution for disaster recovery. In an event of a planned event or unplanned disaster, you can use failover and failback modes as part of your recovery solution. Failover and failback modes help to reduce the time that is required to synchronize remote mirror and copy volumes after you switch between the local (or production) and the intermediate or remote sites during planned and unplanned outages. Although failover transmits no data, it changes the status of a device, and the status of the secondary volume changes to a suspended primary volume. The Failback command transmits data and can go in either direction depending on which device the Failback command is issued to.

Recovery procedures that include failover and failback modes use remote mirror and copy functions, such as Metro Mirror, Global Copy, Global Mirror, Metro/Global Mirror, and FlashCopy.

Note: See the *IBM System Storage DS8000 Command-Line Interface User's Guide* for specific disaster recovery tasks.

Data consistency can be achieved using the following methods:

Manually using external software (without Global Mirror)

If you use Metro Mirror, Global Copy, and FlashCopy functions to create a consistent and restartable copy at your recovery site, you must do a manual and periodic suspend operation at your local site. This means using *freeze and run* commands together with external automated software and then using the FlashCopy function to make a consistent copy of your target volume for backup or recovery purposes. (Automation software is not provided with the storage unit; it must be supplied by the user.)

Note: Freezing of the data is done at the same point-in-time across all links and all storage units.

Automatically (with Global Mirror and FlashCopy)

If you use a two-site Global Mirror or a three-site Metro/Global Mirror configuration, the process to create a consistent and restartable copy at your intermediate or remote site is done using an automated process, with minimal or no interruption to your applications. Global Mirror operations automate the process of continually forming consistency groups. It combines Global Copy and FlashCopy operations to provide consistent data at the remote site. A master storage unit (along with subordinate

storage units) internally manages data consistency using consistency groups within a Global Mirror configuration. Consistency groups can be created many times per hour to increase the currency of data that is captured in the consistency groups at the remote site.

Note: A consistency group is a collection of volumes (grouped in a session) across multiple storage units that are managed together in a session during the creation of 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.

In a two-site Global Mirror configuration, if you have a disaster at your local site and have to start production at your remote site, you can use the consistent point-in-time data from the consistency group at your remote site to recover when the local site is operational.

In a three-site Metro/Global Mirror configuration, if you have a disaster at your local site and you must start production at either your intermediate or remote site, you can use the consistent point-in-time data from the consistency group at your remote site to recover when the local site is operational.

Performing failover and failback operations (without Global Mirror)

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 perform a failover operation 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 accessible.

In a disaster recovery environment, when two storage units are set up in two geographically distinct locations, the storage unit at the production or local site is referred to 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 is moved to Site B during this outage, which makes the target volumes at Site B convert to source volumes and causes them to enter a suspended state. Your original source volumes at Site A remain in the state they were in at the time of the site switch. Table 4 on page 155 provides an example of the implementation of failover and failback operations.

Note: The failover recovery operation does not reverse the direction of a remote mirror and copy pair. It changes 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 continue to update the source volumes that are located at Site A.
- Paths are established from Site A to Site B.
- Volume pairs are in duplex state.

The following steps summarize the actions that you must take to move production to Site B after you initiate a planned outage at Site A.

1. Quiesce applications to cease all write I/O from updating the source volumes when the planned outage window is reached. Quiescing your applications might occur as part of a planned outage, but the delay in processing caused by the quiesce action should be brief.

Note: On some host systems, such as AIX, Windows, and Linux, before you perform FlashCopy operations, you must quiesce the applications that access FlashCopy source volumes. The source volumes must be unmounted (depending on the host operating system) during FlashCopy operations. This ensures that there is no data in the buffers that might be flushed to the target volumes and potentially corrupt them.

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. Create paths in the opposite direction from Site B to Site A depending on your path design and when the source storage unit becomes available. You need the paths in the opposite direction because you want to transfer the updates back to Site A.
4. Rescan (this is dependant on your operating system) your fibre-channel devices. The rescanning removes device objects for the site A volumes and recognizes 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. Initiate a failback recovery operation when your scheduled maintenance is complete. The failback recovery operation initiates the transfer of data back to 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 provides an example of the implementation of failover and failback operations:

Table 4. Failover and failback implementation

Step	Operation	MC connectivity 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	MC connectivity 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 a 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 must 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 must 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 a 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 a 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 discovers which tracks were modified at any site and sends both the tracks that were changed on Site A and the tracks that were 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. The Site B volume is the source volume of the failback operation. This volume was initially the target volume of the relationship.

The following steps summarize the actions that you must take to return production to your local site after initiating a planned outage.

1. Perform a failback operation at Site B. As a result, data that has been modified at Site B for the specified volumes is sent from Site B to Site A. Data that was modified at Site A after the switch to Site B is overwritten by the data from Site B.
2. Quiesce your application (still updating volumes at Site B) to cease all write I/O from updating the source volumes before the application returns to normal operations. 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 Site A to suspended source volumes. The volumes at Site A start the change recording process while in failover mode.
4. Rescan (depending on your operating system) the fibre-channel devices and mount the new source volumes at Site A.
5. Restart the applications in normal production at Site A and check for consistency after the failover operation completes. When your applications are started, all write I/O to the new source volumes are 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)

For this scenario, assume that an unexpected failure occurs at Site A. The failure at Site A causes the volumes to be suspended or causes a mix of suspended and full duplex volume pairs because the input might have been written to those volumes when the failure occurred.

In a disaster recovery environment, when two storage units are set up in two geographically distinct locations, the storage unit at the production or local site is referred to as Site A, and the storage unit at the remote or recovery site is referred to as Site B.

The failover operation is performed on the storage unit that will become the primary. Production is moved to Site B during this outage, which makes the target volumes at Site B convert to source volumes. The volumes are designated as in a suspended state. Your original source volumes at Site A remain in the state that 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.

The following steps summarize the actions you must take to move production to Site B as the result of an unplanned outage and then return production to Site A after it recovers.

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 steps toward the recovery of 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, and 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 performed to restart mirroring either in the reverse direction (remote site to local site) or in the original direction (local site to remote site).

Appendix. Troubleshooting and analyzing problems

The topics in this appendix provide information on troubleshooting and analyzing problems.

Troubleshooting

The topics in this section provide troubleshooting information related to your DS8000. Topics covered include analyzing, verifying, understanding, managing, and handling various problems.

Generic and specific alert traps

SNMP alerts can contain a combination of a generic and a specific alert trap. This list outlines the explanations for each of the possible combinations of generic and specific alert traps.

SNMP alert traps provide information about problems that the storage unit detects. Either you or the service provider must perform corrective action for the following problems, as required:

Generic trap 0 - coldstart

Issued whenever the SNMP agent is re-initializing. Configuration data has changed.

Generic trap 1 - warmstart

Issued whenever the SNMP agent is reinitializing. Configuration data might change.

Generic trap 4 - authenticationFailure

Issued whenever an SNMP message was received but could not be authenticated.

Generic trap 6 - Problem notification and Copy Services traps

The storage unit generates the following specific traps in combination with generic trap 6:

Specific trap 3 - Problem log entry created

This trap displays when a problem log entry is generated in the problem log for the storage unit. The following information is provided for specific trap 3:

```
ssSystemNotifyTrap=yyy/MM/dd hh:mm:ss tzn
Manufacturer=IBM
ReportingMTMS=ttt-mmm*ppzzzzz
ProbNm=nnnnn
LparName=SFaESSbc
FailingEnclosureMTMS=ttt-mmm*ppzzzzz
SRC=sssssssss
EventText=79 char string
Fru1Loc=71 char string
Fru2Loc=71 char string
Fru3Loc=71 char string
```

where *yyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *ttt-mmm* is the machine type and model number; *ppzzzzz* is the serial number (which is a combination of the plant of manufacture and

the sequence number); *nnnnn* is the problem number; *SFaESSbc* is the logical partition name; *sssssssss* is the system reference code (SRC); and *FruNLoc* is the location code of FRU *N*, where *n* is 1, 2, and 3.

Specific trap 100 - Remote mirror and copy links degraded

This trap displays when an established remote mirror and copy path between a primary and secondary logical subsystem pair becomes inoperable, but there is at least one operational path remaining between the two logical subsystems. The following information is provided for specific trap 100:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm LS
PRI: xxx tttt-ooo pp-zzzzz ll
SEC: xxx tttt-ooo pp-zzzzz ll
Path: Type PP PLink SP SLink RC
1: aaaaa bbbb cccccc dddd eeeee gg
.
.
.
w: aaaaa bbbb cccccc dddd eeeee gg
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); *ll* is the LSS number in hexadecimal format; *aaaaa* is "ESCON" or "FIBRE"; *bbbb* is the primary I/O port number; *ccccc* is the primary link address in hexadecimal format; *dddd* is the secondary port number for ESCON direct connection or FIBRE (this value is XXXX if you are connected through an ESCON switch); *eeeeee* is the secondary link address in hexadecimal format; *gg* is the reason code in hexadecimal format (see Table 5) if there is a problem (otherwise, *gg* is "OK" if there is no problem); and *w* represents the last of up to 8 path descriptions. Only established paths are listed.

Table 5 lists remote mirror and copy return codes that can appear for *gg*.

Table 5. Remote mirror and copy return codes

Return code	Return code explanation
X'02'	Initialization failed. The ESCON link rejected threshold was exceeded when attempting to send ELP or RID frames.
X'03'	Timeout. No reason is available.
X'04'	There are no resources available in the primary storage unit for establishing logical paths because the maximum number of logical paths have already been established.
X'05'	There are no resources available in the secondary storage unit for establishing logical paths because the maximum number of logical paths have already been established.
X'06'	There is a secondary storage unit sequence number, or logical subsystem number, mismatch.

Table 5. Remote mirror and copy return codes (continued)

Return code	Return code explanation
X'07'	There is a secondary LSS subsystem identifier (SSID) mismatch, or failure of the I/O that collects the secondary information for validation.
X'08'	The ESCON link is offline. This is caused by the lack of light detection coming from a host, peer, or switch.
X'09'	The establish failed. It is retried until the command succeeds or a remove paths command is run for the path. Note: The attempt-to-establish state persists until the establish path operation succeeds or the remove remote mirror and copy paths command is run for the path.
X'0A'	The primary storage unit port or link cannot be converted to channel mode if a logical path is already established on the port or link. The <i>establish paths</i> operation is not retried within the storage unit.
X'10'	Configuration error. The source of the error is one of the following: <ol style="list-style-type: none"> 1. The specification of the secondary adapter ID does not match the installed ESCON adapter in the primary storage unit. 2. For ESCON paths, the secondary storage unit destination address is zero and the ESCON director (switch) was found in the path. 3. For ESCON paths, the secondary storage unit destination address is not zero and an ESCON director does not exist in the path. The path is a direct connection.
X'14'	The fibre-channel path link is down.
X'15'	The maximum number of fibre-channel path retry operations has been exceeded.
X'16'	The fibre-channel path secondary adapter is not remote mirror and copy capable. This could be caused by one of the following conditions: <ul style="list-style-type: none"> • The secondary adapter is not configured properly or does not have the current firmware installed. • The secondary adapter is already a target of 32 different logical subsystems (LSSs).
X'17'	The secondary adapter fibre-channel path is not available.
X'18'	The maximum number of fibre-channel path primary login attempts has been exceeded.
X'19'	The maximum number of fibre-channel path secondary login attempts has been exceeded.
X'1A'	The primary fibre-channel adapter is not configured properly or does not have the correct firmware level installed.
X'1B'	The fibre-channel path established but degraded due to a high failure rate.
X'1C'	The fibre-channel path was removed due to a high failure rate.

Specific trap 101 - Remote mirror and copy links are inoperable

This trap displays when an established remote mirror and copy path between a primary and secondary logical subsystem pair becomes inoperable, and there are no other operational paths between the two logical subsystems. The following information is provided for specific trap 101:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm LS
PRI: xxx tttt-ooo pp-zzzzz ll
```

```

SEC: xxx tttt-ooo pp-zzzzz ll
Path: Type PP PLink SP SLink RC
1: aaaaa bbbb cccccc dddd eeeee gg
.
.
.
w: aaaaa bbbb cccccc dddd eeeee gg

```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); *ll* is the LSS number in hexadecimal format; *aaaaa* is "ESCON" or "FIBRE"; *bbbb* is the primary I/O port number; *ccccc* is the primary link address in hexadecimal format; *dddd* is the secondary port number for ESCON direct connection or FIBRE (this value is XXXX if you are connected through an ESCON switch); *eeeeee* is the secondary link address in hexadecimal format; *gg* is the reason code in hexadecimal format (see Table 5 on page 160) if there is a problem (otherwise, *gg* is "OK" if there is no problem); and *w* represents the last of up to 8 path descriptions. Only established paths are listed.

Specific trap 102 - Remote mirror and copy links are operational

This trap displays when an established remote mirror and copy path between a primary and secondary logical subsystem pair becomes operational after previously being inoperable. All other established paths between the two logical subsystems are operational. The following information is shown for specific trap 102:

```

ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm LS
PRI: xxx tttt-ooo pp-zzzzz ll
SEC: xxx tttt-ooo pp-zzzzz ll
Path: Type PP PLink SP SLink RC
1: aaaaa bbbb cccccc dddd eeeee gg
.
.
.
w: aaaaa bbbb cccccc dddd eeeee gg

```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); *ll* is the LSS number in hexadecimal format; *aaaaa* is "ESCON" or "FIBRE"; *bbbb* is the primary I/O port number; *ccccc* is the primary link address in hexadecimal format; *dddd* is the secondary port number for ESCON direct connection or FIBRE (this value is XXXX if you are connected through an ESCON switch); *eeeeee* is the secondary link address in hexadecimal format; *gg* is the reason code in hexadecimal format (see Table 5 on page 160) if there is a problem (otherwise, *gg* is "OK" if there is no problem); and *w* represents the last of up to 8 path descriptions. Only established paths are listed.

Specific trap 200 - LSS pair consistency group remote mirror and copy pair error

This trap displays when the first remote mirror and copy pair that is associated with the consistency group enters the suspended state because of an error. This trap is managed by automation code. Reporting for this trap is reenabled when either of the following conditions occurs:

- An extended long busy timeout occurs on the remote mirror and copy pair that is reported in the SNMP trap.
- A consistency group-created operation is directed to the LSS pair that is associated with the consistency group.

The following information is shown for specific trap 200:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm LS LD SR
PRI: xxx tttt-ooo pp-zzzzz ll jj kk
SEC: xxx tttt-ooo pp-zzzzz ll jj
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); *ll* is the LSS number in hexadecimal format; *jj* is the logical volume number in hexadecimal format; and *kk* is the suspension reason code (SRC) in hexadecimal format.

The suspension reason codes explain the cause of the error that suspended the remote mirror and copy group. Table 6 explains the suspension reason codes.

Table 6. Copy Services suspension reason codes

Suspension reason code (SRC)	SRC explanation
X'03'	The host system sent a command to the primary volume of a remote mirror and copy volume pair to suspend copy operations. The host system might have specified either an immediate suspension or a suspension after the copy completed and the volume pair reached a full duplex state.
X'04'	The host system sent a command to suspend the copy operations on the secondary volume. During the suspension, the primary volume of the volume pair can still accept updates but updates are not copied to the secondary volume. The out-of-sync tracks that are created between the volume pair are recorded in the change recording feature of the primary volume.
X'05'	Copy operations between the remote mirror and copy volume pair were suspended by a primary storage unit secondary device status command. This system resource code can only be returned by the secondary volume.
X'06'	Copy operations between the remote mirror and copy volume pair were suspended because of internal conditions in the storage unit. This system resource code can be returned by the control unit of either the primary volume or the secondary volume.

Table 6. Copy Services suspension reason codes (continued)

Suspension reason code (SRC)	SRC explanation
X'07'	Copy operations between the remote mirror and copy volume pair were suspended when the secondary storage unit notified the primary storage unit of a state change transition to simplex state. The specified volume pair between the storage units is no longer in a copy relationship.
X'08'	Copy operations were suspended because the secondary volume became suspended as a result of internal conditions or errors. This system resource code can only be returned by the primary storage unit.
X'09'	The remote mirror and copy volume pair was suspended when the primary or secondary storage unit was rebooted or when the power was restored. Note: The paths to the secondary storage unit might not be disabled if the primary storage unit was turned off. If the secondary storage unit was turned off, the paths between the storage units are restored automatically, if possible. After the paths have been restored, issue the mkpprc command to resynchronize the specified volume pairs. Depending on the state of the volume pairs, you might have to issue the rmpprc command to delete the volume pairs and reissue a mkpprc command to reestablish the volume pairs.
X'0A'	The remote mirror and copy pair was suspended because the host issued a command to freeze the remote mirror and copy group. This system resource code can only be returned if a primary volume was queried.

Specific trap 202 - Primary remote mirror and copy devices on the LSS were suspended because of an error

This trap displays when one or more primary remote mirror and copy devices become suspended because of an error that occurs over a 5-second period. The trap information indicates first pair to suspend as well as all of the primary devices on the primary logical subsystem that were suspended because of the error. The following information is shown for specific trap 202:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm LS LD SR
PRI: xxx tttt-ooo pp-zzzzz ll jj kk
SEC: xxx tttt-ooo pp-zzzzz ll jj
Start: yyyy/mm/dd hh:mm:ss tzn
PRI Dev Flags (1 bit/Dev, 1=Suspended):
ffffffffffff...ffffffffffffffff
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); *ll* is the LSS number in hexadecimal format; *jj* is the logical volume number in hexadecimal format; *kk* is the suspension reason code in hexadecimal format; and *ffffffffffff...ffffffffffffffff* represents 256 1-bit flags, one per primary device on the primary LSS that are shown as a 64-character hexadecimal number. The time indicated on the Start line is the beginning of the error condition period. The time that the alert is sent (as shown on the first line of the trap) indicates the end of the error condition period.

The suspension reason codes explain the cause of the error that suspended the primary remote mirror and copy devices. Table 6 on page 163 explains the suspension reason codes.

Specific trap 210 - Global Mirror initial consistency group successfully formed

This trap displays when the first consistency group in a Global Mirror configuration has formed successfully. The following information is shown for specific trap 210:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
xxx tttt-ooo pp-zzzzz
Session ID: qqqq
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); and *qqqq* is the session identifier.

Specific trap 211 - Global Mirror session is in a fatal state

This trap displays when an error occurs that prevents the formation of a consistency group and that requires you to correct the problem before you can reform the consistency group. The following information is shown for specific trap 211:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
xxx tttt-ooo pp-zzzzz
Session ID: qqqq
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); and *qqqq* is the session identifier.

Specific trap 212 - Global Mirror consistency group failure - Retry will be attempted

This trap displays when a consistency group fails to form. The storage unit automatically continues to attempt to form a consistency group. The following information is shown for specific trap 212:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
xxx tttt-ooo pp-zzzzz
Session ID: qqqq
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); and *qqqq* is the session identifier.

Specific trap 213 - Global Mirror consistency group successful recovery

This trap displays when a consistency group forms successfully after it had previously failed. The following information is shown for specific trap 213:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
xxx tttt-ooo pp-zzzzz
Session ID: qqqq
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); and *qqqq* is the session identifier.

Specific trap 214 - Global Mirror master terminated

This trap displays when a Global Mirror master is terminated by a command. The following information is shown for specific trap 214:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
xxx tttt-ooo pp-zzzzz
Session ID: qqqq
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); and *qqqq* is the session identifier.

Specific trap 215 - Global Mirror FlashCopy at remote site unsuccessful

This trap displays when the FlashCopy operation at the remote site failed in a Global Mirror configuration. The following information is shown for specific trap 215:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
Master: xxx tttt-ooo pp-zzzzz
Slave: xxx tttt-ooo pp-zzzzz
Session ID: qqqq
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); and *qqqq* is the session identifier.

Specific trap 216 - Global Mirror slave termination unsuccessful

This trap displays when a Global Mirror master receives a Terminate command and is unable to terminate a Global Mirror slave. The following information is shown for specific trap 216:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
Master: xxx tttt-ooo pp-zzzzz
Slave: xxx tttt-ooo pp-zzzzz
Session ID: qqqq
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); and *qqqq* is the session identifier.

Specific trap 217 - Global Mirror paused

This trap displays when Global Mirror is paused because of a command received by the Global Mirror master. The following information is shown for specific trap 217:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
xxx tttt-ooo pp-zzzzz
Session ID: qqqq
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); and *qqqq* is the session identifier.

Specific trap 218 - Global Mirror number of consistency group failures exceed threshold

This trap displays when Global Mirror has exceeded the allowed threshold for failed consistency group formation attempts. The following information is shown for specific trap 218:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
xxx tttt-ooo pp-zzzzz
Session ID: qqqq
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *tttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); and *qqqq* is the session identifier.

Specific trap 219 - Global Mirror first successful consistency group after prior failures

This trap displays when Global Mirror has successfully formed a consistency group after one or more formation attempts had previously failed. The following information is shown for specific trap 219:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
xxx tttt-ooo pp-zzzzz
Session ID: qqqq
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *ttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); and *qqqq* is the session identifier.

Specific trap 220 - Global Mirror number of FlashCopy commit failures exceed threshold

This trap displays when Global Mirror has exceeded the allowed threshold of failed FlashCopy commit attempts. The following information is shown for specific trap 220:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
xxx ttt-ooo pp-zzzzz
Session ID: qqqq
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *ttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); and *qqqq* is the session identifier.

Specific trap 221 - Space Efficient Repository or Over-provisioned Volume has reached a warning watermark

This trap is displayed when the Repository has reached the user-defined warning watermark or when physical space is completely exhausted. The following information is shown for specific trap 221:

```
ssEventDescr=yyyy/MM/dd hh:mm:ss tzn
-error_description-
UNIT: Mnf Type-Mod SerialNm
xxx ttt-ooo pp-zzzzz
Volume Type: v
Reason Code: r
Extent Pool ID: ee
Percentage Full: ff
```

where *yyyy/MM/dd hh:mm:ss tzn* is the time that the alert was sent (year, month, day, hour, minute, second, and time zone); *-error_description-* is the event description; *xxx* is the manufacturer; *ttt-ooo* is the machine type and model number; *pp-zzzzz* is the serial number (a combination of the plant of manufacture and the sequence number); *v* is the volume type, over provisioned or repository; *r* is the reason code, 0 = warning watermark and 1 = space exhausted, 100% full; *ee* is the extent pool session identifier in Hex (for example: F2); and *ff* identifies the percentage full

Analyzing service information messages for System z and S/390 hosts

Service information messages (SIMs) are generated by a storage unit for System z and S/390 hosts. Before installation, you can use the customization work sheets to record the SIM severity level and the maximum number of times the storage unit sends the SIMs (0-5 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 resource unit and 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 System Storage DS8000 information:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. Window-Eyes version 5.5 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 System Storage DS8000 information from the keyboard by using the shortcut keys for your browser or screen-reader software. See your browser or screen-reader software Help for a list of shortcut keys that it supports.

Accessing the publications

You can find HTML versions of the IBM System Storage DS8000 information at the following Web site: <http://www.ehone.ibm.com/public/applications/publications/cgibin/pbi.cgi>

You can access the information using screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. Window-Eyes version 5.5 has been tested.

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

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal Use: You may reproduce these Publications for your personal, non commercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these Publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these Publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these Publications, or reproduce, distribute or display these Publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the Publications is detrimental to its interest or, as determined by IBM, the above instructions are not being 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.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Trademarks

This topic lists IBM trademarks.

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

- AIX
- AS/400
- DB2
- developerWorks
- DFSMS/MVS
- DFSMS/VM
- DS4000
- DS6000
- DS8000
- e (logo)
- eServer
- Enterprise Storage Server
- ES/9000
- ESCON
- FICON
- FlashCopy
- GDPS
- Graphically Dispersed Parallel Sysplex
- HACMP
- i5/OS
- iSeries
- IBM
- IntelliStation
- MVS/ESA
- Netfinity
- NetVista
- OpenPower
- Operating System/400
- OS/400
- pSeries
- RS/6000
- Seascape
- SNAP/SHOT
- SP
- S/390
- System/390
- System i
- System i5
- System p
- System p5

- System x
- System z
- System z9
- System z10
- System Storage
- System Storage DS
- Tivoli
- TotalStorage
- Versatile Storage Server
- Virtualization Engine
- VM/ESA
- VSE/ESA
- xSeries
- z/Architecture
- z/OS
- z/VM
- zSeries

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

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

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

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

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

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

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 Union EMC Directive conformance statement

This product is in conformity with the protection requirements of EU Council Directive 2004/108/EC 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.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to CISPR 22/European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

European community contact:

IBM Technical Regulations
Pascalstr. 100, Stuttgart, Germany 70569
Telephone: 0049 (0)711 785 1176
Fax: 0049 (0)711 785 1283
E-mail: tjahn@de.ibm.com

Germany compliance statement

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit.

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 89/336/EWG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung der IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der IBM gesteckt/eingebaut werden.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden: "Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie 89/336/EWG in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) vom 18. September 1998 (bzw. der EMC EG Richtlinie 89/336) für Geräte der Klasse A.

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Konformitätserklärung nach Paragraph 5 des EMVG ist die IBM Deutschland GmbH, 70548 Stuttgart.

Informationen in Hinsicht EMVG Paragraph 4 Abs. (1) 4:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.

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

Taiwan Contact Information

This topic contains the product service contact information for Taiwan.

IBM Taiwan Product Service Contact Information:
IBM Taiwan Corporation
3F, No 7, Song Ren Rd., Taipei Taiwan
Tel: 0800-016-888

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

f2c00790

Glossary

This glossary includes terms for the IBM System Storage 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://w3-03.ibm.com/globalization/page/1728>. 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 6, RAID 10, and 15000 rpm drives. Model 800 supersedes Model F20.
- 1750** The machine type for the IBM System Storage DS6000 series. Models for the DS6000 include the 511 and EX1.
- 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** A hardware machine type for the IBM System Storage DS8000 series. Hardware models for the 2107 include base units 921, 922, 931, 932, 9A2, 9B2 and expansion units 92E and 9AE.
- 2244** A function authorization machine type for the IBM System Storage DS8000 series. The 2244 function authorization machine type corresponds with the 2107 hardware machine type and is used only for purposes of billing and authorizing the licensed functions on the 2107. Function authorization models for the 2244 are related to the type of licensed functions that you order. For example, Model RMC is for the remote mirror and copy function on a 2107 storage unit.
- 239x** Function authorization machine types for the IBM System Storage DS8000 series. These machine types indicate the

warranty period for the licensed functions and they include the following machine types: 2396 (one-year warranty), 2397 (two-year warranty), 2398 (3-year warranty), and 2399 (four-year warranty). Each 239x function authorization machine type corresponds to the 242x hardware machine type that represents the same warranty period. For example, you order a 2398 (3-year warranty) function authorization machine type for a 2423 (3-year warranty) hardware machine. The 239x machine types are used only for purposes of billing and authorizing the licensed functions on the 242x machines. The 239x machine types have one model (Model LFA) with several types of available licenses for that model. For example, Model LFA, RMC license is for the remote mirror and copy function on a 242x storage unit.

242x Hardware machine types for the IBM System Storage DS8000 series. The 242x hardware machine types include machine types 2421 (one-year warranty), 2422 (two-year warranty), 2423 (3-year warranty), and 2424 (four-year warranty). Hardware models for the 242x machine types include base units 931, 932, 9B2 and expansion units 92E and 9AE.

3390 The machine number of an IBM disk storage system. The Enterprise Storage Server, when interfaced to IBM System z 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 computing services or data.
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 System z products, 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 storage unit has four bays, two in each cluster. See also *service boundary*.

bit The smallest unit of computer information, which has two possible states that are represented by the binary digits 0 or 1. See also *byte*.

block A string of data elements recorded or transmitted as a unit. The elements may be characters, words, or physical records. (GC)

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 A string that represents a character and usually consists of eight binary digits that are treated as a unit. A byte is the smallest unit of storage that can be addressed directly. (GC) 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. (GC)

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 storage product and a service provider. The storage product 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. (GC) 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™), pertaining 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*.

central electronics complex

The set of hardware facilities that are associated with a host computer.

channel

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)

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)

A data structure that specifies an I/O operation to the channel subsystem.

channel path

The interconnection between a channel and its associated control units.

channel subsystem

The part of a host computer that manages I/O communication between the program and any attached control units.

channel-subsystem image

In mainframe computing, 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. (GC)

CKD See *count key data*.

CLI See *command-line interface*. See also *IBM System Storage DS™ CLI*.

cluster

1) A partition capable of performing all DS8000 series functions. With two clusters in the DS8000 storage unit, any operational cluster can take over the processing of a failing cluster.

cluster processor complex

The unit within a cluster that provides the management function for the DS8000 series. It consists of cluster processors, cluster memory, and related logic.

command-line interface (CLI)

An interface 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 *IBM System Storage DS CLI*.

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 series 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)

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 series 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 devices, optional features, and program products of the input/output 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 System Storage Management Console*.

contingent allegiance
In mainframe computing, 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. (GC)

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 System z, a storage server with ESCON or OEMI interfaces. The control unit adapts a native device interface to an I/O interface that a System z 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
In mainframe computing, 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 series can emulate 3990-3, TPF, 3990-6, or 2105 control units.

control-unit-initiated reconfiguration (CUIR)
A software mechanism that the DS8000 series uses to request that an operating system of a System z host verify that one or more subsystem resources can be taken offline for service. The DS8000 series 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, that is used for configuring, managing, and monitoring data-copy functions.

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)

In mainframe computing, 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 System z 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*.

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 System z 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 System z 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 System z, 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 System z, 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 System z, 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 series use 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 drive set

A specific number of identical disk drives that have the same physical capacity and rpm.

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.

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 series

See *IBM System Storage DS8000*.

DS8000 Batch Configuration tool

A program that automatically configures a DS8000 storage unit. The configuration is based on data that IBM service personnel enter into the program.

DS Storage Manager

See *IBM System Storage 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*.

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 System z computer peripheral interface. The I/O interface uses System z 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*.

error-recording data set (ERDS)

On System z 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

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 System z 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

System z hosts that attach to the DS8000 series 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 System z, 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 eServer*.

ESSNet

See *IBM TotalStorage Enterprise Storage Server Network*.

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.

FATA See *fibre-channel ATA*.

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 ATA*.

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 ATA (FATA) A hard drive that combines a fibre channel interface with an ATA drive. FATAs, which provide the high performance and capacity of an ATA drive, can be used wherever fibre channel drives can connect.

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. This topology directly connects the storage unit to an open systems host without going through a fabric switch.

Fibre Channel Connection (FICON) A fibre-channel communications protocol that is designed for IBM mainframe computers and peripherals. It connects the storage unit to one or more S/390 hosts using a FICON S/390 channel either directly or through a FICON switch.

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 System z 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*.

Fibre Channel Switched Fabric (FC-SF)

An implementation of the Fibre Channel Standard that connects the storage unit to one or more open systems hosts through a fabric switch or connects one or more S/390 hosts that run LINUX on an Fibre Channel Protocol S/390 channel.

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 series. 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 series 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.

FlashCopy SE

An optional feature of the DS8000 series that allocates storage space on an "as-needed" basis and uses only the required number of tracks to write changed data during the lifetime of the FlashCopy relationship.

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 System z 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 storage unit or on another DS8000 storage unit. 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 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 multiprocessing*.

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

An option on a System i 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 disk assembly*.

HDD See *hard disk drive*.

hdisk An AIX term for storage space.

head 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 storage unit. 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 multiprocessing (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 loop (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, System z) or of the open-systems type, that is connected to the DS8000 series. Hosts are connected through ESCON, FICON, 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 loop*.

HyperPAV (IBM HyperPAV)

An optional licensed function that you can use in conjunction with the parallel access volumes (PAV) function. IBM HyperPAV associates the volumes with either an alias address or a specified base logical volume number. When a host system requests IBM HyperPAV processing and the processing is enabled, aliases on the logical subsystem are placed in an IBM HyperPAV alias access state on all logical paths with a given path group ID. IBM HyperPAV is only supported on FICON channel paths.

i5/OS The IBM operating system that runs the System i family of products.

IBM HyperPAV

See *HyperPAV*.

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 System Storage

The brand name used to identify storage products from IBM, including the IBM System Storage DS8000 series. See also *IBM System Storage DS8000* and *IBM System Storage DS Storage Manager*.

IBM System Storage DS8000

A member of the IBM System Storage Resiliency Family of storage servers and attached storage devices (disk drive modules). The DS8000 series storage product 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 System Storage DS CLI

The command-line interface (CLI) that works with DS8000, DS6000, and 2105 models.

IBM System Storage DS Storage Manager (DS Storage Manager)

Software with a Web-browser interface for configuring the DS8000 series.

IBM System Storage Management Console

(MC) An IBM workstation that acts as the focal point for configuration, Copy Services management, and maintenance for the DS8000 series. It includes a Web browser that provides links to the user interface,

including the DS Storage Manager and the DS8000 Copy Services.

IBM System Storage Multipath Subsystem Device Driver (SDD)

IBM software that provides multipath configuration support for a host system that is attached to storage devices. SDD provides enhanced data availability, dynamic input/output load balancing across multiple paths, and automatic path failover protection.

IBM System Storage Resiliency Family

A set of hardware and software features and products, as well as integrated software and services that are available on the IBM System Storage DS8000 series and the IBM TotalStorage Enterprise Storage Server, Models 750 and 800.

IBM Systems

The IBM brand name for the System i, System p, System x™, and System z families of products.

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.

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 System z 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 System z that enables the system administrator to set priorities for queueing I/Os from different system images. See also *multiple allegiance* and *parallel access volumes*.

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.

IP See *Internet Protocol*.

IPL See *initial program load*.

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*.

licensed machine code

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 series) 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 System z 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 System z, 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. (DS8000 series only)

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 storage unit can have (if CDK only) up to 32 CKD-formatted logical subsystems (8192 CKD logical devices) or (if FBA only) up to 32 fixed-block logical subsystems (8192 fixed-block logical devices). If mixed CKD and FBA, a DS8000 can have up to 16 CKD-formatted logical subsystems (4096 CKD logical devices) and 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 System z control-unit image.

For System z 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 storage unit. 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 System Storage 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 System Storage 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

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*.

MTBF See *mean time between failures*.

Multipath Subsystem Device Driver

See *IBM System Storage 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 System z operating system.

multiplex

The action of transmitting simultaneously.

multiport serial adapter (MSA)

An adapter on the IBM System Storage Management Console that has multiple ports to which a DS8000 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.

near-line

A type of intermediate storage between online storage (which provides constant, rapid access to data) and offline storage (which provides infrequent data access for backup purposes or long-term storage).

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-stripping 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 volumes (PAV)

A licensed 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 System z 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 System z 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 volumes*.

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.

planar The main printed circuit board (PCB) that other PCBs or assemblies plug into. The planar distributes both power and signals and therefore creates a common communications path to whichever device that plugs into it.

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.

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 6

Any form of RAID that can continue to process read and write requests to all of an array's virtual disks in the presence of two concurrent disk failures. See also *device adapter*, *RAID 5*, *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 One or more arrays that are combined to create a logically contiguous storage space.

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 System z, 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 a System z 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 System z, 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 System z 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 System z 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 more storage devices that provides storage capability to a host computer. You can configure more than one storage image on a storage unit. (DS8000 series only)
- 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 i

The IBM family of products which includes the AS/400, eServer iSeries®, eServer i5 and System i5 products.

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.

System p

The IBM family of products which emphasizes performance and includes System p5™, eServer p5, eServer pSeries®, eServer OpenPower™, and RS/6000®.

System x

The IBM family of products which emphasizes industry-standard server scalability and self-managing server technologies. It includes System x3nnn, eServer xSeries, and AMD processor-based eServer servers.

System z

The IBM family of products which emphasizes near-zero downtime and includes System z10, System z9, and eServer zSeries.

System z storage

Storage arrays and logical volumes that are defined in the DS8000 as connected to System z products.

T

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 System z 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*.

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.

System Storage

See *IBM System Storage*.

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 System z, 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 System z, 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*.

Z**z/Architecture**

An IBM architecture for mainframe computers and peripherals. The System z family of products uses the z/Architecture architecture. It is the successor to the S/390 and 9672 family of servers. See also *System z*.

zoning

In fibre-channel environments, the grouping of multiple ports to form a virtual, private, storage network. Ports that are members of a zone can communicate with each other, but are isolated from ports in other zones.

z/OS An operating system for the System z family of products.

z/OS Global Mirror

Remote mirror and copy feature for System z environments. 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*.

Index

Numerics

- 3390 Model A 30
- 524-byte sector, disk formatting 142

A

- accessibility 175
- account setup
 - adding user accounts 93
 - administering user accounts 91
 - unlocking a user account 95
 - unlocking an administrative password 95
- activating licenses 83
- adding
 - users 93
- API, DS Open 13
- arrays
 - across loops 23
- availability features 15

B

- backup and restore commands 147

C

- cascaded relationship
 - Metro/Global Mirror 120
- cascaded volumes 120
- CKD extended address volumes 30
- CLI, command-line interface 12
- command interface
 - DS CLI 118
- commands
 - backup and restore 147
 - cpio 146
 - cplv 143
 - dump and restore 147
 - find 146
 - lslv -l 145
 - lsvg -p rootvg 145
 - migratepv 145
 - migratepv -l 145
 - mklvcopy 144
 - scopy 151
 - xcopy 151
- configuration
 - DS Storage Manager 11
 - reconfiguration 11
- Configuration overview 32
- consistency groups
 - overview 122, 131
 - tunable values
 - consistency group interval 122
 - maximum coordination time 122
 - Maximum time writes are inhibited to the remote site 122
- control unit (CU) 102

- Copy Services
 - 2105 domain, configuring 135
 - across 2105 and 1750 133
 - across 2105 and 2107 133
 - disaster recovery 153
 - licensed functions 29
 - Metro/Global Mirror 130
 - overview 23
 - remote
 - z/OS Global Mirror 130
 - S/390 environment 130
 - System z environment 130
- copying
 - logical volume 143
- cpio commands 146
- cplv command 143
- creating a Metro Mirror using a 2105 137
- CU, control unit 102

D

- data migration
 - associated with physical volumes 145
 - backup and restore 151
 - backup and restore methods 147
 - between storage servers 151
 - copying a logical volume 143
 - direct copy 146
 - from a System i host system 147
 - from System z host systems 148
 - from UNIX-based host systems 149
 - from Windows host systems 151
 - logical volume management software 143
 - migratepv commands 145
 - mirroring 144
 - moving workloads when migrating data between storage units 152
 - offloading workloads when migrating data between storage units 152
 - preparing for 141
 - reformatting existing disk drive modules 142
 - replacing existing storage considerations 142
 - selecting method 141, 142
- DDM
 - formatting 142
- defining
 - the logical control unit 105
- deleting
 - user accounts 94
- description 110
- direct copy method 151
- disabling
 - incremental resynchronization 127
- disaster recovery
 - Copy Services 153
 - failback 154

- disaster recovery (*continued*)
 - failover 154
 - fast restore 116
- disaster recovery solutions
 - GDPS 125
- disk formatting, 524-byte sector 142
- Disk Manager
 - monitoring performance 13
- DS command-line interface 12
- DS Open API 13
- DS storage management
 - Linux graphical install 57
 - Linux graphical upgrade 71
 - Windows uninstall 50
- DS Storage Manager 11
 - accessing 88
 - getting started 82
 - starting 81
- DS8000 interfaces 11
- DS8000 models 1
- DS8000 storage management
 - Linux unattended (silent) mode install 63
 - Linux unattended (silent) mode upgrade 76
 - Windows
 - unattended (silent) mode upgrade 56
 - unattended mode install 43
- DS8000 Storage Manager
 - firewall alert 35
 - operating system support 35
 - Windows graphical install 36
 - Windows graphical upgrade 51
- DS8000 Storage Manager (simulated)
 - hardware requirements 35
 - installation prerequisites 35
- DS8100 series 1
- DS8300 series 1
 - Model 922 1
 - Model 932 1
 - Model 9A2 1
 - Model 9B2 1
- DSFA, Disk Storage Feature
 - Activation 84
- dual hardware management consoles 89
- dual management console 89
- dump and restore commands 147

E

- enabling
 - incremental resynchronization 127
- error messages
 - alert traps 159
 - product-specific 169
 - storage unit 171
- extended address volumes 30
- Extended Remote Copy (XRC) (see z/OS Global Mirror) 130

F

- failback and failover 132
- failover
 - site switch 154, 156
- failover and fallback 153
- fast reverse 116
- Feature Activation (DSFA), IBM Disk Storage 84
- fibre-channel
 - open-systems hosts 16
 - overview 17
 - SAN connections 16
- find commands 146
- FlashCopy
 - changing direction of 109
 - incremental 108
 - multiple relationships 111
 - overview 108
 - persistent overview 111
 - refresh target volume 112

G

- GDPS
 - using GDPS to setup Metro/Global Mirror 125
- generic alert trap 159
- Global Copy
 - overview 114
- Global Mirror 153
 - description 115
 - key terms defined 115
- guidelines
 - for defining the logical control unit 105

H

- hardware management console
 - dual 89
- host attachment
 - overview 16
- host systems
 - System z and S/390 101

I

- IBM Disk Storage Feature Activation (DSFA) 84
- IBM System Storage Resiliency Family
 - Copy Services 107
 - introducing 107
- implementation, RAID 22
- incremental FlashCopy
 - refreshing
 - FlashCopy target 112
- incremental resync
 - three-site Metro/Global Mirror 126
- incremental resynchronization
 - options 127
 - using 127
- information center 10
- informational messages 171
- input/output configuration program (IOCP) 103

- interfaces, DS8000 11
- introducing
 - Metro/Global Mirror 118
- IOCP
 - example 103
 - overview of the IOCP (input/output configuration program) 103

K

- keyboards
 - accessibility features 175

L

- LaunchPad, Windows install 36
- LCU
 - guidelines for defining the LCU (logical control unit) 105
- legal
 - terms and conditions 178
- licensed functions 84
 - Copy Services 29
- licenses
 - Disk Storage Feature Activation (DSFA) 83
 - function authorization documents 84
- limitations
 - remote FlashCopy 138
- Linux
 - graphical install, DS storage management 57
 - graphical upgrade, DS storage management 71
 - unattended (silent) mode install 63
 - unattended (silent) mode upgrade 76
- logical control unit (LCU) 105
- logical partition 97, 98
- logical volume manager
 - software 143
- LPAR 97
 - implementation 97
 - variable 98
- lsiv -l commands 145
- lsvg -p rootvg commands 145

M

- management console
 - changing the password 88
 - dual 89
 - multiple units 87
 - overview 87
- Metro Mirror
 - overview
 - write sequence 113
- Metro/Global Mirror
 - cascaded relationship 120
 - Copy Services 130
 - function 126
 - how it works 120
 - incremental resync 126
 - introducing 118
 - solutions
 - GDPS 125

- Metro/Global Mirror (*continued*)
 - with incremental resynchronization 127
- migratepv -l commands 145
- migratepv commands 145
- migrating data
 - associated with physical volumes 145
 - backup and restore 151
 - backup and restore methods 147
 - between storage servers 151
 - copying a logical volume 143
 - direct copy 146
 - from a System i host system 147
 - from System z host systems 148
 - from UNIX-based host systems 149
 - from Windows host systems 151
 - logical volume management software 143
 - migratepv commands 145
 - mirroring 144
 - moving workloads when migrating data between storage units 152
 - offloading workloads when migrating data between storage units 152
 - preparing for 141
 - reformatting existing disk drive modules 142
 - replacing existing storage considerations 142
 - selecting method 141
- mklvcopy command 144
- Model 921 1, 4
- Model 922 5
- Model 92E 1
- Model 931 1, 4
- Model 932 5
- Model 9A2 5
- Model 9AE 1
- Model 9B2 1, 5
- modifying user accounts 94
- multiple relationship FlashCopy 111

O

- obtaining activation codes 84
- online information center 10
- operations
 - Metro/Global Mirror 120
- optional features codes
 - copy services 101
 - Parallel Access Volumes (PAVs) 101
- overview
 - consistency groups 131
 - FlashCopy 108
 - Global Copy 114
 - host attachment 16
 - input/output configuration program (IOCP) 103
 - logical partition 97
 - Remote Mirror and Copy 112
 - sessions management 118

P

- panel
 - rack operator 21
- parallel access volume (PAV)
 - understanding static and dynamic 102
- parallel access volumes
 - requirements 102
- passwords
 - changing the management console password 88
 - use criteria 93
- path
 - connections 129
- PAV (parallel access volumes) 102
- PAVs
 - defining an LCU 105
 - increasing the formatted 3390s 106
 - IOCP program 103
 - reconfiguring an LCU 106
- performance gathering
 - Disk Manager 13
- performing
 - site switch
 - failover 154, 156
- persistent FlashCopy
 - overview 111
- power
 - unit emergency power off (UEPO) switch 20
- powering a storage unit on and off 20
- problem states 169
- product-specific error messages 169

R

- rack operator panel 21
- RAID
 - implementation 22
 - RAID 10 overview 22
 - RAID 5 overview 22
 - RAID 6 overview 22
 - RAID overview 22
- RAID 10
 - arrays across loops 23
- reconfiguring an LCU 106
- refresh volume, FlashCopy 108
- remote FlashCopy limitations 138
- Remote Mirror and Copy 112
- removing
 - user accounts 94
- Replication
 - copy services functions 14
- requirements
 - parallel access volume 101
- return codes 159
- reversing
 - FlashCopy relationship 109

S

- S/390 host
 - ESCON attachment overview 17
 - FICON attachment overview 18
- S/390 hosts
 - parallel access volume 101

SAN

- connections with fibre-channel adapters 16
- scenarios
 - configuring
 - Copy Services 2105 domain 135
 - scopy commands 151
- SCSI-FCP 17
- security
 - defining user groups 91
- service information messages (SIMs) 169
- sessions, understanding 118
- SNMP alert traps 159
- SNMP messages 159
- SNMP trap
 - messages 159
- specific alert trap 159
- SRC 159
- SRN 159
- storage area network
 - connections with fibre-channel adapters 16
- storage system logical partition
 - overview 97
- storage unit
 - error messages 171
 - powering on and off 20
- switch
 - unit emergency power off (UEPO) 20
- synchronous data mirroring 113
- System z host
 - ESCON attachment overview 17
 - FICON attachment overview 18
- System z hosts
 - parallel access volume 101

T

- TotalStorage Productivity Center
 - copy services 14
 - for Disk 13
 - Replication 14
- Trademarks 179
- troubleshooting
 - informational messages 171
 - normal operation problems 171
 - problem states 169
 - product-specific error messages 169
 - service information messages (SIMs) 169
 - SNMP trap messages 159
 - storage unit error messages 171

U

- UEPO 20
- uninstall DS storage management
 - Windows 50
- unit emergency power off switch (UEPO) 20
- updating the I/O configuration 105
- user administration
 - adding user accounts 93
 - administering user accounts 91
 - deleting user accounts 94
 - modifying user accounts 94

- user administration (*continued*)
 - unlocking a user account 95
 - unlocking an administrative password 95
- users
 - adding 93
 - defining access 91
 - deleting user accounts 94
 - modifying user accounts 94
 - removing user accounts 94
 - user groups 91
- using FlashCopy
 - with remote mirror and copy 110

V

- variable LPAR 98

W

- Windows
 - DS Storage Manager upgrade,
 - graphical 51
 - graphical install, DS8000 Storage Manager 36
 - install, LaunchPad 36
 - uninstall, DS storage management 50
 - Windows unattended (silent) mode install 43
 - Windows unattended (silent) mode upgrade 56

X

- xcopy commands 151
- XRC (Extended Remote Copy) (see z/OS Global Mirror) 130

Z

- z/OS Global Mirror 130

Readers' comments — we would like to hear from you

**IBM System Storage DS8000
User's Guide
Version 1 Release 4**

Publication No. SC26-7915-02

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

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. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Send your comments to the address on the reverse side of this form.

If you would like a response from IBM, please fill in the following information:

Name

Address

Company or Organization

Phone No.

E-mail address



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



Printed in USA

SC26-7915-02

