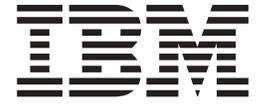


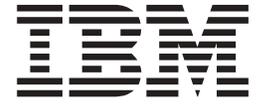
IBM System Storage DS6000



Installation, Troubleshooting, and Recovery Guide

Version 1 Release 2

IBM System Storage DS6000



Installation, Troubleshooting, and Recovery Guide

Version 1 Release 2

Note:

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

First Edition (September 2006)

This edition replaces GC26-7678-08 and all previous versions of GC26-7678.

This edition also applies to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2004, 2006. All rights reserved.**

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

Contents

About this guide vii

Notices and publication information . . ix

Safety notices	ix
Environmental notices	ix
Product recycling and disposal	ix
Battery return program	x
Conventions used in this guide	xi
Related information	xii
DS6000 series library	xii
Other IBM publications	xiii
Ordering IBM publications	xvii
Web sites	xviii
How to send your comments	xix

Summary of Changes for GC26-7925-00 IBM System Storage DS6000 Installation, Troubleshooting, and Recovery Guide xxi

Chapter 1. Introduction 1

Where to go for information	1
Installing the DS6000 and creating an initial storage configuration overview	2
Introduction to the IBM System Storage DS6000 series	2
Overview of the DS6000 series models	3
DS6800 (Models 1750-511 and 1750-522)	3
DS6000 expansion enclosure (Models 1750-EX1 and 1750-EX2)	5
Performance features	6
RAID implementation	7
RAID 5 overview	7
RAID 10 overview	7
Licensed functions	8
DS6000 Interfaces	9
IBM System Storage DS Storage Manager	9
The DS command-line interface	9
DS open application programming interface	10
IBM TotalStorage Productivity Center for Disk	11
IBM TotalStorage Productivity Center for Replication	11
Host systems attachment overview	12
IBM System Storage Resiliency Family	12
Disaster recovery using Copy Services	17
Comparison of licensed functions	18
Logical configuration overview	19
Simulated configuration overview	22
Real-time configuration overview	22
Express configuration overview	22

Chapter 2. Preparing your system for installation 25

Preinstallation planning and verification	25
---	----

Service-clearance and floor-load requirements	25
Power outlet requirements	25
Input voltage requirements	26
Power connector requirements	26
Planning for network and communications requirements	27
Gathering installation information	29

Chapter 3. Installing the DS6000 series hardware 33

Verifying the ship group	33
Assembling tools and equipment	34
Understanding ESD requirements	34
Positioning the rack	35
Removing the hardware resources	35
Removing the front display panel	37
Removing the disk drive module	37
Removing the power supply	38
Removing the rear display panel	39
Removing the battery backup units	40
Removing the processor cards	41
Installing the support rails	42
Installing the server enclosure in the rack	44
Installing expansion enclosures in the rack	46
Replacing the hardware resources	47
Replacing the disk drive module	47
Replacing the power supply	48
Replacing the rear display panel	49
Replacing the battery backup unit	50
Replacing the processor cards	51
Installing the front display panel	52
Installing the host systems and I/O adapters	53
Routing the cables	53
DS6000 fabric zoning	54
Connecting network cables	55
Installing the fiber optic cables and SFPs	56
Connecting expansion enclosures	56
Connecting hosts to server enclosure processor cards	61
Connecting power cords	63
Powering on	64
Verifying status through LEDs	65
Setting the IP address	66
Enabling multiple IP addresses on the DS6000 Storage Manager	68
Installing a modem on the management console	69

Chapter 4. Installing the DS6000 Storage Manager 73

Internet browser support	74
Management console requirements	74
Installing the DS Storage Manager on the Windows operating system	76

Installing the DS Storage Manager on the Windows operating system using the graphical mode	76
Installing the IBM System Storage DS Storage Manager on the Windows operating system in unattended (silent) mode	82
Starting the DS Storage Manager on a Windows operating system	89
Setting up the DS Storage Manager	90
Creating a configuration file (simulated only)	91
Creating a storage unit (simulated only)	92
Creating a storage complex (simulated only)	93
Locating the serial number, model, and signature	93
Assigning a storage unit to a storage complex (real-time only)	94
Verifying the current code level.	95
Activating licensed functions	95
Defining multiple management consoles (real-time only)	97
Specifying storage unit day and time (real-time only)	98
Reviewing storage unit network settings (real-time only)	99
Defining customer contacts (real-time only).	99
Setting up Call Home, SNMP, and SIM notifications (real-time only)	100
Configuring I/O ports	101
Creating host systems	102

Chapter 5. Installing the DS6000 Command-Line Interface 105

Preparing to install the DS CLI	105
Preparing for the DS CLI installation	105
Preparing to install the DS CLI on an OpenVMS System	108
Installing the DS CLI	109
Installing the DS CLI using the unattended (silent) mode	109
Installing the DS CLI using the console mode	113
Installing the DS CLI using the graphical mode	119
Correcting the JVM-not-found error	125
Installing the DS CLI on your OpenVMS system	126
Completing DS CLI postinstallation	132
Using the DS CLI application	132
DS CLI command help	136
Obtaining the serial (storage image ID) number using the DS CLI	138
Activating your machine and feature licenses using the DS CLI	139
Setting up user accounts using the DS CLI	140
Configure support options using the DS CLI	141
Creating a default configuration setup with a profile file	147
OpenVMS system integration	152
Enhancing the command console LUN for DS CLI use	152
Enhancing the OpenVMS system messages	153
Enabling OpenVMS to use the DS CLI help	154
Java Runtime Environment considerations for DS CLI	154
Quota considerations for DS CLI	155

Chapter 6. Enabling support options 157	
Registering for the My Support function	158
Activating the Call Home function using the DS CLI.	159
Activating remote support (real-time only)	160
Unconfiguring modem phone numbers (real-time only)	160

Chapter 7. Creating an initial logical storage configuration using DS Storage Manager 163

Creating an express configuration	163
Creating open systems volumes using express configuration	163
Creating zSeries volumes using express configuration	165
Creating iSeries volumes using express configuration	166
Creating a custom logical storage configuration	167
Creating arrays.	168
Creating ranks	169
Creating extent pools.	170
Creating zSeries LCUs	171
Creating zSeries volumes	172
Creating open systems volumes	173
Creating open systems volume groups	174
Applying a configuration (simulated only)	175
Enabling host system I/O	176

Chapter 8. Creating a logical storage configuration using the DS command-line interface 177

Configuring new fixed block storage using the DS CLI.	177
Creating extent pools for fixed block volumes using the DS CLI	177
Creating arrays for fixed block volumes using the DS CLI	179
Creating a rank using the DS CLI	180
Correcting a rank-related configuration error	181
Creating fixed block volumes using the DS CLI	182
Correcting a fixed block configuration error	184
Creating fixed block volume groups using the DS CLI	184
Configuring fibre-channel I/O ports using the DS CLI	186
Creating SCSI host port connections using DS CLI.	187
Configuring new count key data storage using DS CLI.	189
Creating count key data extent pools using the DS CLI	189
Creating arrays for CKD volumes using the DS CLI.	191
Creating a rank for CKD volumes using the DS CLI.	192
Creating logical control units for CKD volumes using DS CLI	193
Creating count key data volumes using the DS CLI.	194

About this guide

This publication provides installation, troubleshooting, and recovery information related to your DS6000™.

All of the information in this guide is also contained in the IBM System Storage™ DS6000 Information Center. You can access the Information Center from any screen of the DS Storage Manager by clicking the help icon.

DS6000 Information Center



The most up-to-date version of the Information Center is available at:

<http://publib.boulder.ibm.com/infocenter/ds6000ic/index.jsp>

Note: A link for this Web site is available on the Start page of the IBM System Storage DS6000 Customer Information CD that is shipped with your DS6000.

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.

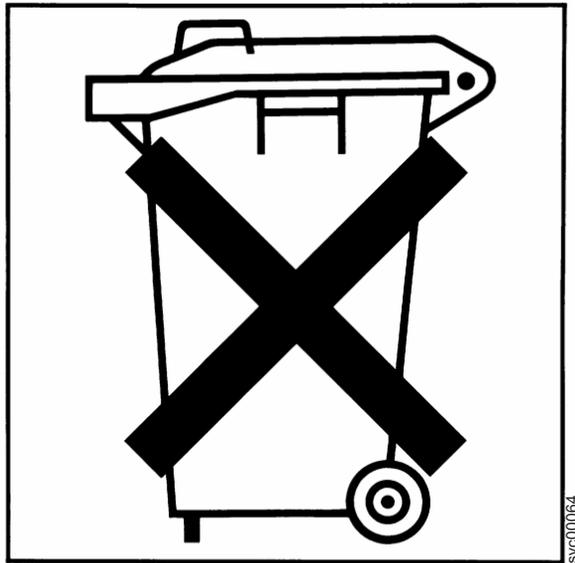
In accordance with the European WEEE Directive, electrical and electronic equipment (EEE) is to be collected separately and to be reused, recycled, or recovered at end of life. Users of EEE with the WEEE marking per Annex IV of the WEEE Directive, as shown above, must not dispose of end of life EEE as 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 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/batteryrecycle.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

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

- The publications that make up the IBM® System Storage™ DS6000 series library
- Other IBM publications that relate to the DS6000 series
- Non-IBM publications that relate to the DS6000 series

See “Ordering IBM publications” on page xvii for information about how to order publications in the IBM System Storage DS6000 series publication library. See “How to send your comments” on page xix for information about how to send comments about the publications.

DS6000 series library

These customer publications make up the DS6000 series library.

Unless otherwise noted, these publications are available in Adobe portable document format (PDF) on a compact disc (CD) that comes with the storage unit. If you need additional copies of this CD, the order number is SK2T-8825. These publications are also available as PDF files by clicking on the **Documentation link** on the following Web site:

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

See “Ordering IBM publications” on page xvii 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 DS6000 configuration and Copy Services relationships. The CLI application provides a set of commands that you can use to write customized scripts for a host system. The scripts initiate predefined tasks in a Copy Services server application. You can use the CLI commands to indirectly control Remote Mirror and Copy and FlashCopy® configuration tasks within a Copy Services server group.	GC26-7681 (See Note.)
<i>IBM System Storage DS6000: Host Systems Attachment Guide</i>	This guide provides guidelines for attaching the DS6000 to your host system and for migrating to fibre-channel attachment from a small computer system interface.	GC26-7680 (See Note.)
<i>IBM System Storage DS6000: Introduction and Planning Guide</i>	This guide introduces the DS6000 product and lists the features you can order. It also provides guidelines for planning the installation and configuration of the storage unit.	GC26-7679
<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 DS6000. SDD creates single paths (vpaths) 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.	SC30-4096
<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.	GC35-0493

Title	Description	Order Number
<i>IBM System Storage DS6000 Messages Reference</i>	This publication provides explanations of error, information, and warning messages that are issued from the DS6000 user interfaces.	GC26-7682
<i>IBM System Storage DS6000 Installation, Troubleshooting, and Recovery Guide</i>	This publication provides reference information for installing and troubleshooting the DS6000. It also discusses disaster recovery using Copy Services.	GC26-7678
<i>IBM System Storage DS6000 Quick Start Card</i>	This is a quick start guide for use in installing and configuring the DS6000 series.	GC26-7685
Note: No hardcopy book is produced for this publication. However, a PDF file is available from the following Web site: http://www-1.ibm.com/servers/storage/support/disk		

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 xvii for information about ordering these and other IBM publications.

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

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

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

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

Title	Description	Order Number
<i>IBM OS/390 Hardware Configuration Definition User's Guide</i>	This guide explains how to use the Hardware Configuration Data application to perform the following tasks: <ul style="list-style-type: none"> • Define new hardware configurations • View and modify existing hardware configurations • Activate configurations • Query supported hardware • Maintain input/output definition files (IODFs) • Compare two IODFs or compare an IODF with an actual configuration • Print reports of configurations • Create graphical reports of a configuration • Migrate existing configuration data 	SC28-1848
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide</i>	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08.	SC26-7350
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide</i>	This guide describes the IBM SAN Fibre Channel Switch and the IBM TotalStorage ESS Specialist. It provides information about the commands and how to manage the switch with Telnet and the Simple Network Management Protocol (SNMP).	SC26-7349
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide</i>	This publication describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S16. It is intended for trained service representatives and service providers.	SC26-7352
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide</i>	This guide introduces the IBM SAN Fibre Channel Switch 2109 Model S16 and tells you how to manage and monitor the switch using zoning and how to manage the switch remotely.	SC26-7351
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the Enterprise Storage Server. It provides you with a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment.	SG24-6113
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>IBM TotalStorage Solutions Handbook</i>	This handbook, from the IBM International Technical Support Organization, helps you understand what makes up enterprise storage management. The concepts include the key technologies that you must know and the IBM subsystems, software, and solutions that are available today. It also provides guidelines for implementing various enterprise storage administration tasks so that you can establish your own enterprise storage management environment.	SG24-5250

Ordering IBM publications

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 DS6000 series and other IBM storage products.

Type of Storage Information	Web Site
Concurrent Copy for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
Copy Services command-line interface (CLI)	http://www-1.ibm.com/servers/storage/support/software/cscli/
DS6000 Information Center	http://publib.boulder.ibm.com/infocenter/ds6000ic/index.jsp
DS6000 series publications	http://www-1.ibm.com/servers/storage/support/disk Click Documentation .
FlashCopy for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
Host system models, operating systems, and adapters that the storage unit supports	http://www.ibm.com/servers/storage/disk/ds6000/interop.html Click Interoperability matrix .
IBM Disk Storage Feature Activation (DSFA)	http://www.ibm.com/storage/dsfa
IBM storage products	http://www.storage.ibm.com/
IBM System Storage DS6000 series	http://www-1.ibm.com/servers/storage/disk/ds6000
IBM version of the Java (JRE) that is often required for IBM products	http://www-106.ibm.com/developerworks/java/jdk/
Multiple Device Manager (MDM)	http://www.ibm.com/servers/storage/support/ Click Storage Virtualization .
Remote Mirror and Copy (formerly PPRC) for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
SAN fibre channel switches	http://www.ibm.com/storage/fcswitch/
Storage Area Network Gateway and Router	http://www-1.ibm.com/servers/storage/support/san/
Subsystem Device Driver (SDD)	http://www-03.ibm.com/servers/storage/support/software/sdd
Technical notes and product tips	http://www.ibm.com/servers/storage/support/disk/ds6800/ Click Technical notes on the Troubleshooting tab.
z/OS Global Mirror (formerly XRC) for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this information or any other DS6000 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 GC26-7925-00 IBM System Storage DS6000 Installation, Troubleshooting, and Recovery 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

This edition contains the following new information:

- The following new sections have been added to Chapter 9:
 - Checking for code updates
 - Downloading code updates
 - Installing code updates
- A new section, "Powering off," has been added to Chapter 11.

Changed Information

This edition contains the following changed information:

- The name of Chapter 9 has been changed to "Upgrading the DS6000 Code Level."
- Information on upgrading the code level has been moved to Chapter 9, and the code load procedures have been revised.
- The "Powering on" section has been revised.
- A modem kit is now included in the ship group, and installing the modem on the management console is now a required hardware installation step. The procedure for installing the modem has been moved to Chapter 3, "Installing the DS6000 series hardware."

Deleted Information

The following sections have been deleted from this edition:

- "Interenclosure Power Operations." Some of the information from this section has been revised and moved to the Powering on section, and some of the information from this section has been revised and moved to the new "Powering off" section.

Chapter 1. Introduction

The *IBM System Storage Installation, Troubleshooting, and Recovery Guide* contains the information that you need to begin using your IBM System Storage DS6000.

The following information is contained within this book:

- Hardware installation instructions.
- Software installation instructions for both the IBM System Storage DS Storage Manager and the IBM System Storage DS CLI.
- Logical storage configuration instructions. You can perform logical storage configuration using either the IBM System Storage DS Storage Manager or the IBM System Storage DS CLI.
- Troubleshooting information for when you are unable to access the IBM System Storage DS6000 Information Center.
- Recovery information for when you are unable to access the IBM System Storage DS6000 Information Center.

Where to go for information

The IBM System Storage DS6000 Information Center contains all of the user assistance information for the DS6000 series.

The user assistance information includes, but is not limited to, field upgrades and maintenance procedures, reference information, conceptual overviews, planning and installation procedures, and configuration scenarios. You can access the Information Center using any one of the following steps:

- Click **Start, Programs, DS Storage Manager, and IBM System Storage DS Information Center** on your Windows operating system where the IBM System Storage DS Storage Manager has been installed.
- Click **Start here for configuration instructions** or **Using the Information Center** from the Welcome page of the IBM System Storage DS Storage Manager.
- Click **IBM System Storage DS6000 Information Center** on the start page of the IBM System Storage DS6000 series Customer Information CD that is shipped with your DS6000.
- Click the question mark icon (Launch Information Center) that is located in the top right corner on any page in the IBM System Storage DS Storage Manager.
- Visit the IBM Publications Center, click **Information Centers and Libraries** in the navigation on the left, and click the link for the IBM System Storage DS6000 Information Center in the list of available Information Centers.
- Go directly to the IBM System Storage DS6000 Information Center at <http://publib.boulder.ibm.com/infocenter/ds6000ic/index.jsp>.

Note: A link for this Web site is available on the start page of the IBM System Storage DS6000 Customer Information CD that is shipped with your DS6000. The online version of the Information Center may not correspond to the version installed on the management console. Use the installed version when possible.

For up-to-date product tips, click **Technical notes** on the Troubleshooting tab of the following Support for System Storage DS6800 Web site:

Installing the DS6000 and creating an initial storage configuration overview

This section contains the high-level process that you must complete to install and configure your IBM System Storage DS6000 series.

The following steps outline, at a very high level, the tasks that you must complete to install and initially configure your DS6000 series. Each step below involves its own set of steps.

1. Prepare your system for installation and familiarize yourself with the safety requirements. See “Preinstallation planning and verification” on page 25 for more information, including a table that lists the information that you must gather before you begin your installation.
2. Install the hardware. See Chapter 3, “Installing the DS6000 series hardware,” on page 33 for detailed installation steps.
3. Install and set up the DS Storage Manager. See Chapter 4, “Installing the DS6000 Storage Manager,” on page 73 for information about installing and setting up the DS Storage Manager. This section also includes steps to activate your licensed functions, which you must do before you can logically configure your storage unit, and steps to set up and test the Call Home feature, which must be set up to automatically report errors to IBM support personnel.
4. (Optional) Install the DS CLI. See Chapter 5, “Installing the DS6000 Command-Line Interface,” on page 105 for more information and detailed installation steps.
5. Set up support capabilities and options. See Chapter 6, “Enabling support options,” on page 157 for more information about setting up and using your support options, such as SMTP (Call Home), SNMP, and zSeries service information messages.
6. Create the initial logical configuration using either the DS Storage Manager or the DS CLI. See Chapter 7, “Creating an initial logical storage configuration using DS Storage Manager,” on page 163 for information about configuring your storage using the DS Storage Manager. See Chapter 8, “Creating a logical storage configuration using the DS command-line interface,” on page 177 for information about configuring your storage using the DS CLI.

Introduction to the IBM System Storage DS6000 series

The IBM System Storage DS6000 series is a member of the family of DS products and is built upon 2 Gbps fibre-channel technology that provides RAID-protected storage with advanced function, scalability, and increased addressing capabilities.

The DS6000 series is an enterprise-class storage solution that offers high reliability and performance in a modular midrange package. The DS6000 series provides storage sharing and consolidation for a wide variety of operating systems and mixed server environments.

The DS6000 series offers high scalability and excellent performance. With the DS6800 (Models 1750-511 and 1750-522), you can install up to 16 disk drive modules (DDMs). The minimum storage capability with 4 DDMs is 292 GB. The DS6000 series also offers 500 GB, 7 200 rpm Fibre-channel ATA (FATA) disk drives, so the maximum storage capability with 16 DDMs for the DS6800 model is 8 TB. FATA drives combine a fibre-channel interface with an ATA drive, which provides

additional capacity to address specific application and business requirements. You can add FATA disk drives to the DS6000 series to satisfy the needs of various fixed-content, data-archival, reference data, and near-line applications that require large amounts of storage capacity at a lower cost.

If you want to connect more than 16 disks, you can use the optional DS6000 expansion enclosures (Models 1750-EX1 and 1750-EX2). A storage system supports up to 128 disk drives for a total of up to 64 TB of storage.

The DS6800 measures 5.25-in. high and is available in a 19-in. rack mountable package with an optional modular expansion enclosure of the same size to add capacity to help meet your growing business needs.

The DS6000 series addresses business efficiency needs through its heterogeneous connectivity, high performance, and manageability functions, thereby helping to reduce total cost of ownership.

The DS6000 series offers the following major features:

- PowerPC 750GX processors
- Dual active controllers that provide continuous operations through the use of two processors that form a pair that back up each other
- A selection of 2 GB fibre-channel (FC) disk drives, including 73 GB, 146 GB, and 300 GB sizes with speeds of 10 000 or 15 000 revolutions per minute (rpm)
- 500 GB, 7 200 rpm FATA disk drives
- 2 GB fibre-channel and FICON host attachments of up to 8 ports, which can be configured with an intermix of Fibre-channel Protocol (FCP) and FICON
- Fibre-channel Arbitrated Loop (FC-AL) switched device attachment of up to 2 dual loops
- Storage virtualization
- Battery backed mirrored cache
- Fully redundant power and cooling system
- Disaster recovery and Copy Services features

You can also view the DS6000 e-learning overview from the DS6000 Storage Manager Welcome page at <http://www-1.ibm.com/support/docview.wss?rs=1112&uid=s8g1S7001165>. The e-learning overview provides an animated presentation about installation and configuration, service and support, and management tasks.

Note: Some DS6000 series features and functions might not be available or be supported in all environments. You can find current information on supported environments, prerequisites, and minimum operating systems levels at the following Web site: <http://www.ibm.com/servers/storage/disk/ds6000/interop.html>

Overview of the DS6000 series models

The DS6000 series offers a base enclosure model with storage and optional expansion enclosures.

DS6800 (Models 1750-511 and 1750-522)

The DS6000 offers the following features:

- Two fibre channel processor cards
- PowerPC 750GX 1 GHz processor

- 4 GB of cache
- Two battery backup units (one per each processor card)
- Two ac/dc power supplies with imbedded enclosure cooling units
- Eight 2 Gb/sec device ports
- Connectivity with the availability of two to eight fibre-channel/FICON host ports. The host ports autonegotiate to either 2 Gbps or 1 Gbps link speeds.
- Attachment to 7 DS6000 expansion enclosures.



The DS6800 is a self-contained 3U enclosure that can be mounted in a standard 19-inch rack. The DS6800 comes with authorization for up to 16 internal fibre-channel DDMs or FATA drives, offering up to 8 TB of storage capability. The DS6800 allows up to 7 DS6000 expansion enclosures to be attached. A storage system supports up to 128 disk drives for a total of up to 64 TB of storage.

Note: The intermix of drive types is not allowed within an enclosure.

The DS6800 system offers connectivity with the availability of two to eight fibre-channel/FICON host ports. The 2 GB fibre-channel/FICON host ports, which are offered in longwave and shortwave, autonegotiate to either 2 Gbps or 1 Gbps link speeds. This flexibility supports the ability to exploit the potential benefits offered by higher performance, 2 Gbps SAN-based solutions, while also maintaining compatibility with existing 1 Gbps infrastructures. In addition, with the maximum of eight host ports enabled, the DS6800 system can be configured with an intermix of Fibre Channel Protocol (FCP) and FICON. This can help protect your investment in fibre-channel adapters and increase your ability to migrate to new servers.

The DS6800 system offers connectivity support across a broad range of server environments, including IBM eServer™, zSeries®, iSeries™, and pSeries® servers as well as servers from Sun Microsystems, Hewlett-Packard, and other Intel-based providers. The ability to easily partition the capacity of the DS6800 system storage among the attached environments supports storage consolidation requirements in dynamic and changing environments.

DS6000 expansion enclosure (Models 1750-EX1 and 1750-EX2)

The DS6000 series expansion enclosure contains the following features:

- Two expansion processor cards. Each processor card provides the following:
 - 2 inbound ports (2 Gb/sec.)
 - 2 outbound ports (2 Gb/sec.)
 - 1 fibre-channel switch per processor card
- Expansion disk enclosure that holds up to 16 fibre-channel DDMs
- Two ac/dc power supplies with imbedded enclosure cooling units
- Supports attachment to DS6800



The 3U DS6000 expansion enclosure can be mounted in a standard 19-inch rack. The front of the enclosure contains the docking sites where you can install up to 16 DDMs.

The DDMs are installed in a horizontal position with a locking handle. The rear of the enclosure provides the docking sites for the power supplies and the processor cards.

You can attach the DS6800 and expansion enclosure by using the processor card interfaces at the rear of the enclosure. A system display panel is also located at the rear of the enclosure.

Performance features

The DS6000 series is built upon 2 Gbps fibre-channel technology that can help bring high availability RAID-protected storage with scalable capacity, increased addressing capabilities, and connectivity to a wide range of storage area network (SAN) applications.

The DS6000 series provides the following technology and hardware to meet today's on demand business environments:

Integrated RAID controller technology

The DS6000 series features IBM's 32-bit PowerPC microprocessor, a fourth generation processing technology.

High availability

The DS6000 series is designed with component redundancy to eliminate single points of hardware failure, and no single point of repair other than the enclosure.

Industry standard fibre-channel disk drives

The DS6000 series offers a selection of 2 GB fibre-channel disk drives, including 300 GB drives, and 500 GB, 7200 rpm fibre-channel ATA (FATA) disk drives, which allow the DS6000 series to scale up to a capacity of 64 TB.

Parallel Access Volumes

Parallel Access Volumes (PAV) can provide significant performance enhancements in zSeries environments by enabling simultaneous processing for 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 can be automatically managed to help the workload meet its performance objectives and reduce overall queuing. PAV is an optional feature for the DS6000 series and also requires the purchase of the FICON Attachment feature.

LUN and volume management

LUN and volume creation and deletion is nondisruptive. When you delete a LUN or volume, the capacity can immediately be reused. You can configure LUN and volumes to span arrays, which allows larger LUNs and volumes.

Addressing capabilities

The DS6000 series allows the following:

- Up to 32 logical subsystems
- Up to 8192 logical volumes
- Up to 1040 volume groups

Simplified storage management for zSeries with z/OS

The DS6000 series supports a new 65 520 cylinder 3390 volume. This volume option has a capacity of approximately 55.7 GB. It helps relieve addressing constraints, improve disk resource use, and improve storage administrator productivity by providing the ability to consolidate multiple disk volumes into a single address.

System management

The DS6000 series provides online and offline configuration capability features and a graphical user interface (GUI) that is designed to offer increased ease of use.

A single command line interface (CLI) supports both logical configuration and copy services.

Online Information Center

The online Information Center is an information database that provides the opportunity to quickly familiarize yourself with the major aspects of the DS6000 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.

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.

With RAID implementation, the DS6000 series offers 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 for the DS6000 series can be configured as RAID 5, RAID 10, or a combination of both. 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.

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 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 10 overview

RAID 10 provides high availability by combining features of RAID 0 and RAID 1. The 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.

Licensed functions

Activate licensed features before you configure your DS6000 series.

Before you can configure your DS6000 series, you must activate your licensed features to enable the functions purchased on your machine. The DS6000 series is licensed at the following levels:

- **Machine licensing** uses licensed machine code to activate base functions on your machine. When you receive a DS6800 and/or DS6000 expansion enclosure, you receive an licensed machine code agreement. The use of the machine constitutes acceptance of the license terms outlined in the licensed machine code agreement. The DS6800 system requires licensed machine code level 5.0.0, or later.

Some DS6000 series features and functions may not be available or supported in all environments. Current information on supported environments, prerequisites, and minimum operating systems levels is available at: <http://www.ibm.com/servers/storage/disk/ds6000/interop.html>

- **Operating environment licensing** manages the machine operating environment and is required on every DS6800 system. The extent of IBM authorization acquired through the DS6800 feature numbers (50xx) must cover the physical capacity of the DS6800 system, where system is defined as the base enclosure and all attached expansion enclosures.

If the operating environment license has not been acquired and activated on the machine, disk drives installed within the DS6800 system cannot be logically configured for use. Upon activation, disk drives can be logically configured up to the extent of authorization.

As additional disk drives are installed, the extent of IBM authorization must be increased by acquiring additional DS6800 feature numbers (5xxx). Otherwise, the additional disk drives cannot be logically configured for use.

- **Feature licensing** controls the licenses of features of each DS6800. Each DS6800 licensed function feature number enables the use of, and establishes the extent of, IBM authorization for licensed functions that are acquired for a DS6800 system.

Each licensed function feature number is applicable only for the specific DS6800 (by serial number) for which it was acquired and is not transferable to another serial numbered DS6800.

To activate the feature licenses for your DS6000 series, you must access the Disk Storage Feature Activation (DSFA) application from the IBM Web site.

DS6000 Interfaces

This section describes the following interfaces:

- IBM System Storage DS Storage Manager
- IBM System Storage DS Command-Line Interface
- IBM System Storage DS Open Application Programming Interface
- IBM TotalStorage® Productivity Center for Disk
- IBM TotalStorage Productivity Center for Replication

IBM System Storage DS Storage Manager

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

The DS Storage Manager software must be installed on a computer (management console) that you provide.

You can install the DS Storage Manager as a GUI (graphical mode) or as an unattended (silent) mode installation for the supported operating systems. 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 issue a command from a command prompt window. The response file is a template on the IBM System Storage Storage Manager CD. The DS Storage Manager can be accessed from any location that has network access using a Web browser. It offers you the following choices:

Simulated configuration

This component provides the ability to create or modify logical configurations when your storage unit is disconnected from the network. After creating the configuration, you can save it and then apply it to a network attached storage unit at a later time.

This simulated configuration tool is installed on your server and can be used for the configuration of a DS6000 series at initial installation.

Real-time configuration

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

Express configuration

This component provides the simplest and fastest configuration method.

Copy services

This component allows you to process copy services functions.

The DS command-line interface

The IBM System Storage DS command-line interface (CLI) enables open systems hosts to invoke and manage FlashCopy and Metro and Global Mirror functions through batch processes and scripts.

The DS CLI provides a full-function set of commands that allows you to check your storage unit configuration and to perform specific application functions, when necessary. Many of the values that are used or reported with the DS CLI commands are hexadecimal.

Note: Before you can use the DS CLI commands, ensure that you have met the following conditions:

- Your management console must be equipped with the DS Storage Manager graphical user interface (GUI).
- The GUI must be installed as a full management console installation management type.
- Your storage unit must be configured (part of DS Storage Manager postinstallation instructions).
- You must activate your licensed functions (part of DS Storage Manager postinstallation instructions) before you can use the CLI commands that are associated with Copy Services functions.
- You must not start more than 100 DS CLI sessions simultaneously. Starting more than 100 DS CLI sessions simultaneously can result in connection problems.
- You did not install the DS CLI on a Windows 64-bit operating system.

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

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

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 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, 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 TotalStorage 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

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.

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>

Host systems attachment overview

The DS6000 storage unit provides a variety of host attachments so that you can consolidate storage capacity and workloads for open-systems hosts, S/390® hosts, and eServer™ zSeries® hosts.

You can configure the storage unit for fibre-channel adapters, for support of Fibre Channel Protocol (FCP) and fibre connection (FICON®) protocol .

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

- Port number
- Worldwide port name

You can configure switch ports that are attached to the storage unit in more than one zone. This enables multiple host system adapters to share access to the storage unit fibre-channel ports. Shared access to a storage unit fibre-channel port might come from host platforms that support a combination of bus adapter types and the operating systems.

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 provide disaster recovery, data migration, and data duplication functions.

As a part of the IBM System Storage Resiliency family, the DS6000 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
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. For zSeries environments, you can also use 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 DS6000 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.

The point-in-time and remote mirror and copy features are supported across various server environments such as zSeries, pSeries[®], iSeries[™] and 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. For zSeries 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 the point-in-time feature code 52xx (which can include codes 5200, 5201, 5202, 5203, or 5204).

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 allows for data backup and disaster recovery. To use it, you must purchase the remote mirror and copy code 53xx (which can include codes 5300, 5301, 5302, 5303, or 5304). DS6000 storage units can participate in remote mirror and copy solutions with the DS8000 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 z/Series 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.)

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 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 DS6000 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 configuration, the process to create a consistent and restartable copy at your 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.

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

Table 1. 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.

Table 1. Comparison of licensed functions (continued)

Licensed function	Description	Advantages	Considerations
Global Copy	Continuous copy without data consistency	Nearly unlimited distance, suitable for data migration, only limited by network and channel extenders capabilities.	Copy is normally fuzzy but can be made consistent through synchronization.
Global Mirror	Asynchronous copy	Nearly unlimited distance, scalable, and low 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.

Logical configuration overview

Before you configure your DS6000 series, 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 four disks form an array site. Logical groupings of one or two array sites 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 with IBM eServer zSeries or S/390 systems, 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 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 1 on page 21 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
| will be 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.

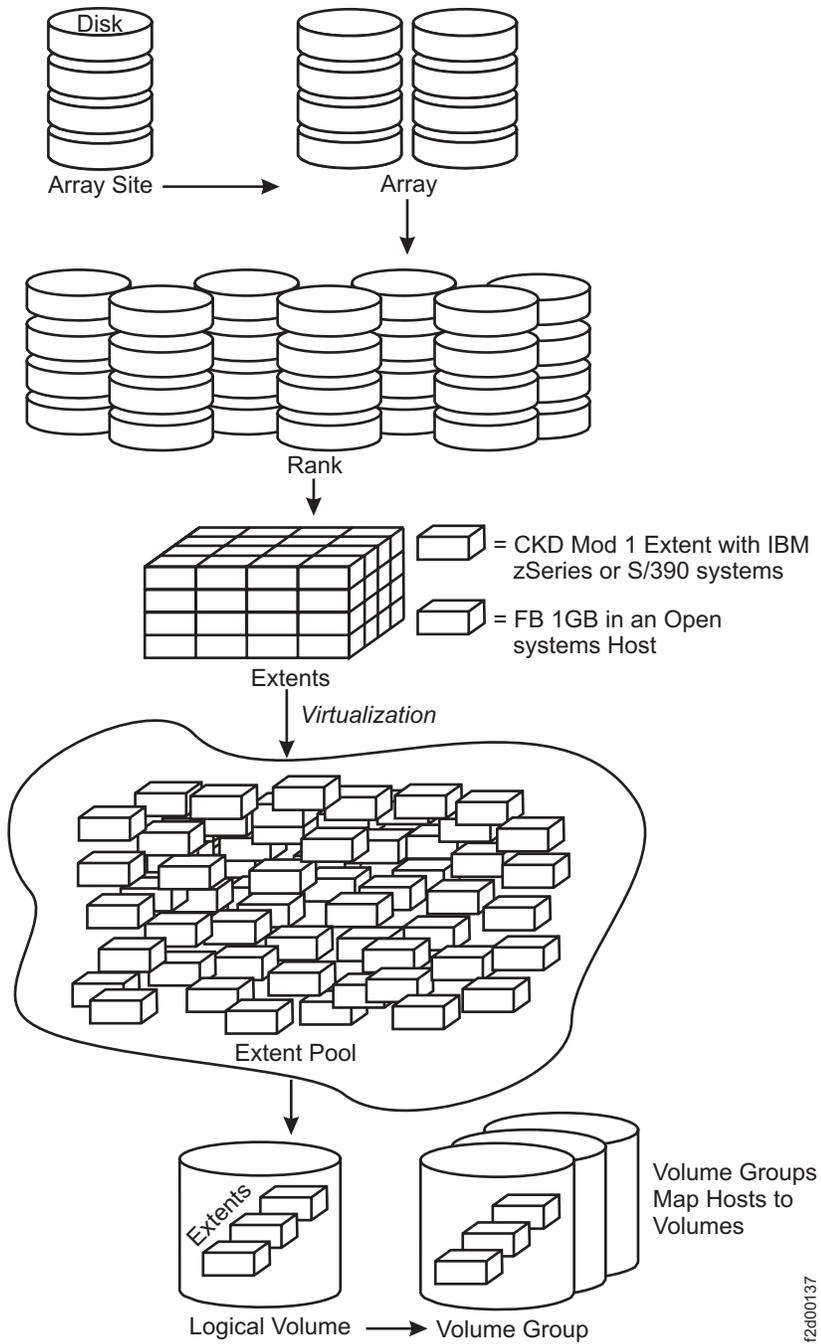


Figure 1. Logical configuration sequence

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.

The custom configuration option is ideal if you want to customize your configuration. The Express Configuration option shortens and simplifies the configuration process, but only allows you to create one host. This is ideal if you

want only one host and you want to quickly and easily set up and configure your storage because most storage-appropriation decisions have already been made.

Simulated configuration overview

You can use the simulated (offline) configuration method to create or import a new simulated instance of your DS6000 series.

Use the simulated configuration method to perform the following tasks:

- Import a physical or a logical configuration, or both, from an existing storage server.
- Apply logical configurations to a new or fully deconfigured storage server.
- Collect communication settings for the storage complex and the management system.
- Apply communication settings as part of the logical configuration process.
- From a single interface, work with a new storage complex or storage unit, and view existing storage complexes and storage units.
- Create, save, and open configuration documents for later reference and retention purposes.
- Print configuration reports.
- Export configuration data in a spreadsheet ready format.

Real-time configuration overview

You can use the real-time (online) configuration method to manage physical and logical configurations from existing storage complexes, storage servers, and storage units over your network.

Use the real-time configuration method to perform the following tasks:

- Construct and apply valid logical configuration actions on new or fully deconfigured storage servers at the time that each action is initiated.
- Complete and apply valid logical configuration actions on existing storage complexes, storage servers, and storage units at the time that each action is initiated.
- Collect communication settings for the storage complex and the storage manager.
- Apply communication settings as part of the logical configuration process.

Express configuration overview

Express configuration provides the simplest and fastest method to configure a storage complex.

Some configuration methods require extensive time. Because there are many complex functions that are available to you, you are required to make several decisions during the configuration process. However, with the express configuration method, the storage server makes several of those decisions for you, decreasing the number of steps that are required to configure volumes and hosts. This eliminates extensive configuration process time and simplifies the task for you.

The express configuration method is ideal for the following users:

- Novice users with little knowledge of storage concepts who want to quickly and easily set up and begin using storage

- Expert users who want to quickly configure a storage complex by allowing the storage server to make decisions for the best storage appropriation

Using the express configuration method, you can perform the following tasks:

- Configure open systems, iSeries, and zSeries volumes
- Create a volume group
- Create a host
- Map a volume group to a host attachment

Chapter 2. Preparing your system for installation

This section provides information that you need to prepare the DS6000 for installation into a rack.

Preinstallation planning and verification

Use this information to prepare the DS6000 for installation into a rack.

Prior to installation, ensure that you complete the following steps to prepare the DS6000 for installation into a rack. See the *IBM System Storage DS6000 Introduction and Planning Guide* for more detailed information.

1. Prepare the site to meet all safety and environmental requirements.
2. Prepare the site to meet all space and floor load requirements.
3. Prepare the site to meet all power requirements.
4. Prepare the site to meet all network and communications requirements.
5. Plan your storage complex setup with the DS6000 customization worksheets. These worksheets are available in the *IBM System Storage DS6000 Introduction and Planning Guide*.

Service-clearance and floor-load requirements

For each DS6000 series model or expansion unit, service personnel must be able to open the front and rear covers to perform service.

Service clearance refers to the empty space in front of, next to, or behind a storage unit that a service person needs in order to access the unit. The service clearance of adjacent units can overlap. Weight distribution areas that are required to handle floor loading do not overlap.

Use the following minimum service clearances as you plan the placement of your system.

- At the front of the unit, allow a minimum of 121.9 cm (48-in.)
- At the rear of the unit, allow a minimum of 76.2 cm (30-in.)
- At each side of the unit, allow a minimum of 5.1 cm (2-in.)

Power outlet requirements

Plan for the required power outlets when planning for the installation of your storage units.

The following power outlets are required:

- Two independent power outlets for the two DS6000 power line cords needed by each base model and expansion model.

Important: To eliminate a single point of failure, the outlets must be independent. This means that each outlet must use a separate power source and each power source must have its own wall circuit breaker.

- Two outlets that are within 3.1 m (10 ft.) of the external management console. Typically, these outlets will be in a rack that you provide.

Input voltage requirements

Specific input voltages are required for the DS6000 series.

Table 2 lists the input voltages and frequencies that the DS6000 series power line cords support. The values apply to both the primary line cords to any storage or expansion enclosure in a DS6000 series. DS6000 series power inputs are single phase.

Table 2. DS6000 series input voltage requirements

Characteristic	Value
Nominal input voltages	100-127 RMS V ac 200-240 RMS V ac
Minimum input voltage	90 RMS V ac
Maximum input voltage	264 RMS V ac
Input frequencies	50 ± 3.0 Hz 60 ± 3.0 Hz

Power connector requirements

Ensure that the site where you plan to install the DS6000 storage units meets the power connector requirements.

Table 3 provides the power line cords and the inline connectors and receptacles types that they support. Find the power line cord row that is appropriate for your site and ensure that the site where you plan to install the DS6000 storage units meets the power connector requirements that are shown in that row.

Phase rotation on 3-phase power-cable connectors is counterclockwise as you look at the power line cord plug. Phase rotation is clockwise as you look at the face of the power receptacle at your installation site.

Attention:

- Low-voltage, three-phase installations (200 - 240 V) require wall circuit breakers that have a rating of 50 to 60 A.
- High-voltage, three-phase installations (380 - 480 V) require wall circuit breakers that have a rating of 30 to 35 A.

Do not exceed the wire rating of the facility.

Table 3. DS6000 power line cords

Power Line Cord Feature Code	Power Line Cord Description	Inline Connector	Receptacle
1090 (See Notes 1, 2, 3, and 5.)	Line cord for U.S., Canada, LA, and AP: Three phase 50 or 60 Hz 60 amperes	7428-78	7324-78
1091 (See Note 1, 2, and 5.)	Line cord for EMEA: Three phase 50 Hz 35 amperes	Not applicable	Hard-wired
1092 (See Notes 1, 2, 4, 5, and 7.)	Line cord for Japan, U.S., and Canada: Three phase 50 or 60 Hz 60 amperes	HBL460C9W	HBL460R9W

Table 3. DS6000 power line cords (continued)

Power Line Cord Feature Code	Power Line Cord Description	Inline Connector	Receptacle
1093 (See Notes 1, 2, 3, and 6.)	Line cords for Chicago: Three phase 50 or 60 Hz 60 amperes	7428-78	7324-78
<p>Notes:</p> <ol style="list-style-type: none"> All power line cords are rated at 600 V ac, except line cords for Japan, which are 250 V ac. Each cord has 3 poles and 4 wires. The conductor size for non-EMEA and Japan line cords is 6 awg. The conductor size for EMEA cords is 6 mm². Connectors and receptacles are supplied by Russellstoll. Preferred connectors and receptacles are supplied by Hubbell. With the exception of the Chicago line cord, all line cords extend 4.2 m (14 ft.) from lowest point where they exit the frame to the mating face of the plug or bare leads. Chicago line cords extend 1.8 m (6 ft.) from lowest point where they exit the frame to the mating face of the plug connector or bare leads. The IEC309 receptacle must be installed in a metal-backed box with the green wire ground-connected to the grounding lug within the box. Ensure continuity between the box and the metallic shielding of the liquid-tight conduit. 			

Amperage ratings of the power connector for the expansion unit must match the ratings that are used in the base model. For example, because a Model uses a 60 A connector, any Model expansion model that is attached to that base model must use a 60 A connector.

Planning for network and communications requirements

You must locate your DS6000 series in a location that meets the network and communications requirements.

Keep in mind the following network and communications issues when you plan the location and interoperability of your storage complex:

- Management console network configuration
- Host attachment requirements
- Remote support connection requirements
- SAN considerations
- Dial-up modem connections

Host attachment communication requirements

This list describes requirements and other considerations for connecting host attachments in your network.

- You must use worldwide port names to uniquely identify fibre-channel adapter cards that are installed in your host system.
- For open-system hosts with fibre-channel adapters, keep in mind that fibre channel architecture provides a variety of communication protocols. Each interconnected storage unit within the architecture is referred to as a *node*, and each host is also a node. Each node corresponds to one or more ports. (In the case of fibre-channel I/O adapters, these ports are fibre-channel ports.) Each port attaches to a serial-transmission medium that provides duplex communication

with the node at the other end of the medium. You can configure your network structure based on one of three basic interconnection topologies (network structures):

- Point-to-point
- Switched fabric
- Arbitrated loop

See the *IBM System Storage DS6000 Host Systems Attachment Guide* for more information about these supported topologies.

- The maximum distance between a host fibre-channel port and the following network components is 300 meters (984 ft) with a shortwave adapter and 10 km (6.2 miles) with a longwave adapter.
 - Fabric switches
 - Fabric hubs
 - Link extenders
 - Storage unit fibre-channel port

The maximum distance might be greater than 10 km (6.2 miles) when a link extender provides target initiator functions or controller emulation functions.

Note: Do not use link extenders with emulation functions on links over which Remote Mirror and Copy operations are performed. This is because of the additional path delay that is introduced by these units.

- Because the fibre channel architecture allows any fibre-channel initiator to access any fibre-channel device, without access restrictions, this can represent a security exposure. You must set the fibre-channel access modes to the proper setting. See the *IBM System Storage DS6000 Host Systems Attachment Guide*, for more information about fibre-channel access modes.

SAN requirements and considerations

These requirements and considerations can help you plan for a DS6000 series that attaches to a SAN.

A fibre-channel storage area network (SAN) is a specialized, high-speed network that attaches servers and storage devices. With a SAN, you can perform an any-to-any connection across the network using interconnect elements such as routers, gateways, hubs, and switches.

When you connect your DS6000 storage units to a SAN, you must meet the following requirements:

- Fibre-channel I/O adapters must be configured to operate in a point-to-point mode fabric topology. See the *IBM System Storage DS6000 Host Systems Attachment Guide* for more information.

Also keep the following considerations in mind:

- Fibre-channel SANs can provide the capability to interconnect open systems and storage in the same network as S/390 and zSeries host systems and storage.
- A single fibre-channel I/O adapter can have physical access to multiple fibre-channel ports on the storage unit.

Remote support and Call Home connection requirements

You must meet these requirements if you will use remote support or the Call Home feature.

You must provide an outside connection, such as one of the following:

- A virtual private network (VPN) connection over your local area network or over a dial-up connection through a modem that is connected to the management console. If you initiate a VPN connection over a modem, you must have an analog phone line that is available for the modem to connect to.
- An Internet connection through your firewall that allows IBM to connect to your storage management system.

Gathering installation information

Use the table in this section to gather the information that you will need to complete the installation of your IBM System Storage DS6000.

Table 4 provides a place for you to gather vital information for your DS6000 installation. You will enter this information while you install and configure your DS6000. Not all of the items that are listed in this table are required for all types of installations. Your data requirements might differ slightly from those that are listed here.

Note: You might have already gathered some of this information while reviewing the *IBM System Storage DS6000 Introduction and Planning Guide*. If so, you can reuse that information here.

Important: When you configure the Call Home feature, you must provide complete, accurate, up-to-date contact information, including the correct telephone number. Incomplete or incorrect contact information can cause a delay in IBM's ability to respond to a Call Home event. If your contact information changes after you initially configure Call Home (for example, if you move the DS6000 to a new location) be sure to update your customer contact information.

Each processor card must have access to the SMTP server to allow the Call Home feature to work properly, but the management console does not require access to the SMTP server.

Table 4. Preparing installation information

What you need	Value
Management console IP address	
IP address for processor card 1	
IP address for processor card 2	
Network mask for processor card 1	
Network mask for processor card 2	
DNS server for processor card 1	
DNS server for processor card 2	
Gateway address for processor card 1	
Gateway address for processor card 2	

Table 4. Preparing installation information (continued)

What you need	Value
Operating Environment Licenses	
Machine type	
Machine model number	
Machine serial number	
Primary DNS for the management console	
Secondary DNS for the management console	
Customer business name	
Customer service account number	
Country	
Country code (for telephone number)	
Area or city code (for telephone number)	
Primary telephone number	
Primary telephone number extension	
Alternate telephone number	
Alternate telephone number extension	
Mailing address	
Building, floor, room location	
City	
State or province	
Zip code or postal code	
Primary country code	
Primary area or city code	
Alternate country code	
Alternate area or city code	
SMTP Call Home server host name	
SMTP Call Home server IP address	
SMTP Call Home port numbers	
SNMP Call Home server host name	
SNMP Call Home server IP address	
SNMP (Call Home) destination port number	

Table 4. Preparing installation information (continued)

What you need	Value
SNMP (Call Home) community name	
SNMP (Call Home) system contact name	
Host system worldwide port names	

Chapter 3. Installing the DS6000 series hardware

This section directs you through the steps of installing the hardware for your IBM System Storage DS6000 series enclosures.

Follow these steps to install the hardware for the server or expansion enclosure:

1. Verify the ship group. You can find additional steps in “Verifying the ship group.”
2. Assemble your tools and equipment. You can find additional steps in “Assembling tools and equipment” on page 34.
3. Position the rack. You can find additional steps in “Positioning the rack” on page 35.
4. Remove the hardware resources. You can find additional steps in “Removing the hardware resources” on page 35.
5. Install the support rails. You can find additional steps in “Installing the support rails” on page 42.
6. Install the server enclosure in the rack. You can find additional steps in “Installing the server enclosure in the rack” on page 44.
7. Install the expansion enclosures in the rack. You can find additional steps in “Installing expansion enclosures in the rack” on page 46.
8. Replace the hardware resources. You can find additional steps in “Replacing the hardware resources” on page 47.
9. Install the host systems and I/O adapters. You can find additional information in “Installing the host systems and I/O adapters” on page 53.
10. Connect all cables. You can find additional steps in “Routing the cables” on page 53.
11. Power on the enclosures. You can find additional steps in “Powering on” on page 64.
12. Verify the status of the enclosure. You can find additional information in “Verifying status through LEDs” on page 65.
13. Set the IP addresses of the processor cards. You can find additional steps in “Setting the IP address” on page 66.
14. Install the modem on the management console. You can find additional steps in “Installing a modem on the management console” on page 69.

Verifying the ship group

Complete this task to verify the ship group.

1. Move the DS6000 to the site.
2. Remove the DS6000 from its shipping container and check the contents.
Ensure that the DS6000 has shipped with the following standard ship group items:
 - The enclosure, including:
 - Two server enclosure processor cards or two expansion enclosure processor cards
 - Two power supplies/fan assemblies

- Two battery backup units (server enclosure only) or two battery blanks (expansion enclosure only)
 - 16 blank trays (your server enclosure might come with up to 16 disk drive modules in place of blank trays)
 - One service information card tray (installed in the rear of the server enclosure, which is located below the lower processor card)
 - Rack-mounting hardware kit, including:
 - Two rails (right and left assembly)
 - Two M5 flathead screws (installed in the rail assembly)
 - Four M5 hex screws
 - Cables, including:
 - Two standard power cords
 - Two 1-meter inline power cords
 - One Ethernet crossover cable (server enclosure only)
 - One serial conversion cable (server enclosure only)
 - Two 13-meter Ethernet cables (server enclosure only)
 - Modem kit
 - Software, including:
 - Microcode CD
 - CLI CD
 - Management Console software CD (SDD is included)
 - License Machine Code Agreement
 - License Function Authorization document
 - Statement of Limited Warranty
 - Code Reference Flyer
 - Electrostatic discharge (ESD) wrist strap
3. Ensure that the items that are listed in your packing slip match what is in the box and that any optional items that you ordered are included in the list. For example, if you ordered fibre-channel cables, Small-form Factor Pluggables (SFPs), or optional power cords, ensure that you received them. If any items are missing or damaged, contact IBM customer support before proceeding.

Assembling tools and equipment

Complete this task to assemble tools and equipment.

1. Assemble the tools and equipment that you will need for installation. These might include:
 - A flat-head screwdriver
 - A holding cart on which to place the hardware resources.
2. When you are ready, begin to prepare the rack.

Understanding ESD requirements

Always wear an electrostatic discharge (ESD) wrist strap that is properly connected to the ESD ground bracket when you service this machine. This prevents possible damage to the hardware and decreases any possible impact to your operations.

Purpose

To prevent damage when you work with ESD-sensitive parts, perform these instructions carefully.

- Never touch or work on any electronic circuits without wearing the ESD wrist strap.
- Keep the ESD-sensitive part in a special ESD bag until you are ready to install the part into the machine.
- Make the fewest possible movements with your body to prevent an increase of static electricity from clothing, fibers, carpets, and furniture.
- If you are instructed to do so, switch off the machine power before you remove ESD-sensitive parts.
- Just before you touch the ESD-sensitive part, discharge to the machine any static electricity in your body by touching the metal frame or the cover of the machine. If possible, keep one hand on the frame when you install or remove an ESD-sensitive part.
- Do not place any ESD-sensitive parts on the machine cover or on a metal table because large metal objects can become discharge paths if they are not grounded. If you must set aside an ESD-sensitive part, first place it into the special ESD bag.
- Prevent ESD-sensitive parts from being accidentally touched by others.
- Be very careful when you work with ESD-sensitive parts in cold weather. Low humidity and heating increase static electricity.

Positioning the rack

Before you install the DS6000 in a rack, keep in mind the following considerations:

- Review the safety considerations.
 - Install the DS6000 in a recommended 10 - 40°C (50 - 104°F) environment.
 - To ensure proper airflow, do not block the front or rear of the rack.
 - To ensure rack stability, load the rack starting at the bottom.
 - If you install multiple components in the rack, do not overload the power outlets.
 - Always connect the server or expansion enclosure to a properly grounded outlet.
 - It is recommended that the rack power be connected to at least two different power circuits or sources. Connecting the rack power to at least two different power circuits or sources allows the enclosure to continue to operate if one of the power sources fails.
1. Move, unpack, and level the rack at the installation site (if needed).
 2. Remove the external rack panels.
 3. Install any additional interface cables and power cables.

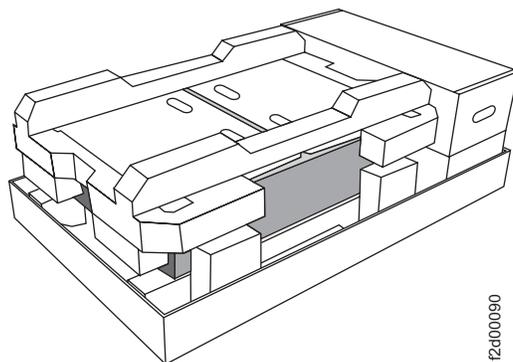
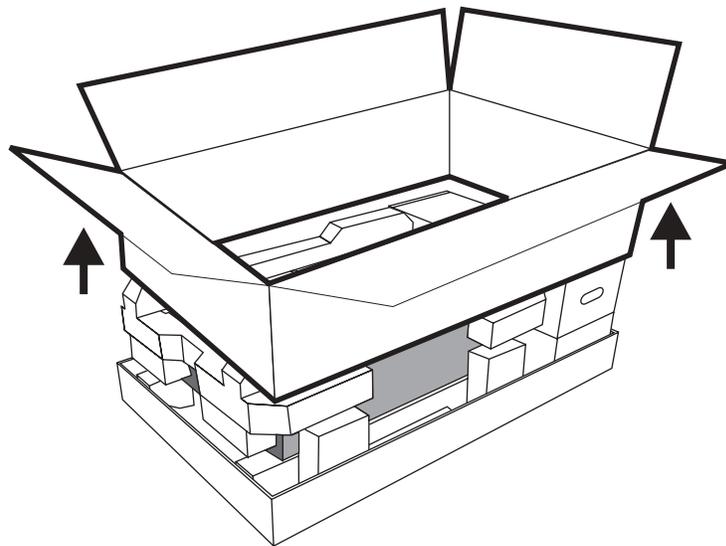
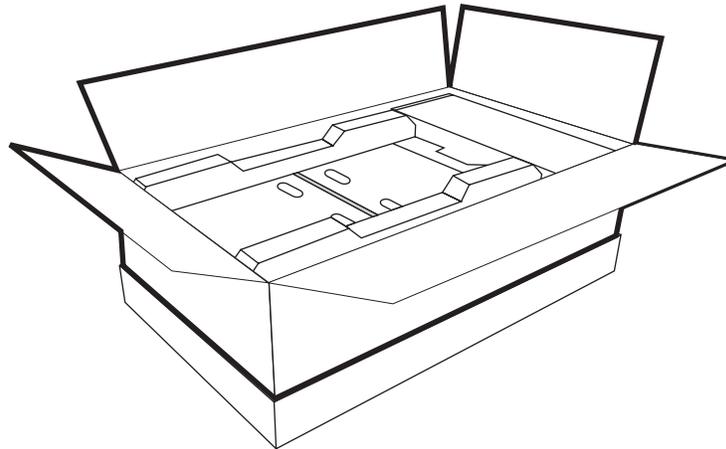
Removing the hardware resources

Complete this task to remove the hardware resources from the enclosure to minimize the weight of the DS6000 before you install it in the rack.

The fully populated enclosure weighs more than 109 lbs (49.5 kg). You need at least three people to lift and install a fully populated DS6000 into the rack. If you do not have three people that are available to lift the DS6000 into the rack, you must remove the hardware resources from the enclosure to reduce the weight

before you lift it. If you have three people to lift and install the DS6000, you might not find it necessary to remove the resources before installation. If this is the case, you can skip this section.

The following images show how to remove the cardboard sides of the shipping box from around the enclosure.



12d00090

1. Make sure to follow ESD procedures by using an anti-static wrist strap. You will also need a cart or level surface to hold the resources.

2. Remove all the resources in the DS6000 storage and server enclosures to minimize the weight.

Removing the front display panel

Complete this task to remove the front display panel from an enclosure.

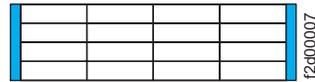


Use approved ESD procedures to prevent damage.

System alerts and LED indicators cannot display properly while the front display panel is removed or nonfunctional.

Important: When you remove the front display panel from an enclosure, ensure that all resources from a particular server or expansion enclosure are reinstalled into the same physical enclosure and in the same location inside that enclosure.

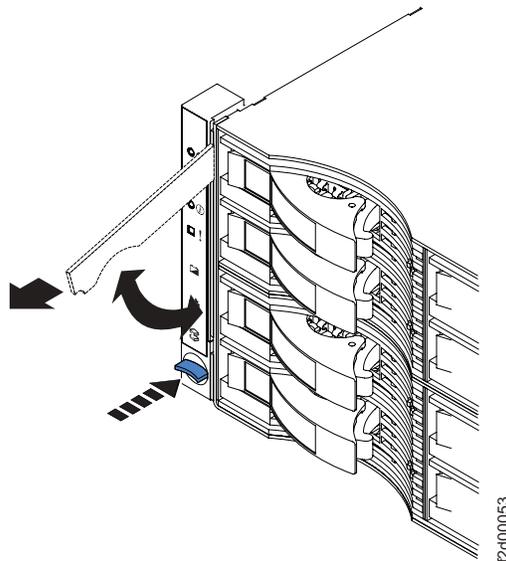
The front display panels are located on the left and right of the front of the enclosure.



To remove the front panel display, perform the following steps:

1. Press the blue release button. This releases the handle.
2. Pull the handle up. This action partially ejects the bottom of the display panel.
3. Use the handle to pull the bottom of the panel up and release the top portion of the panel. This completely removes the display panel.

This image show a front display panel removal procedure.

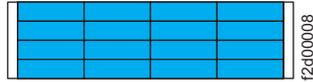


Removing the disk drive module

Complete this task to remove the disk drive modules from the enclosure to reduce the weight of the enclosure during installation.

 Use approved ESD procedures to prevent damage.

The disk drive modules are located in rows and columns on the front of the storage unit.

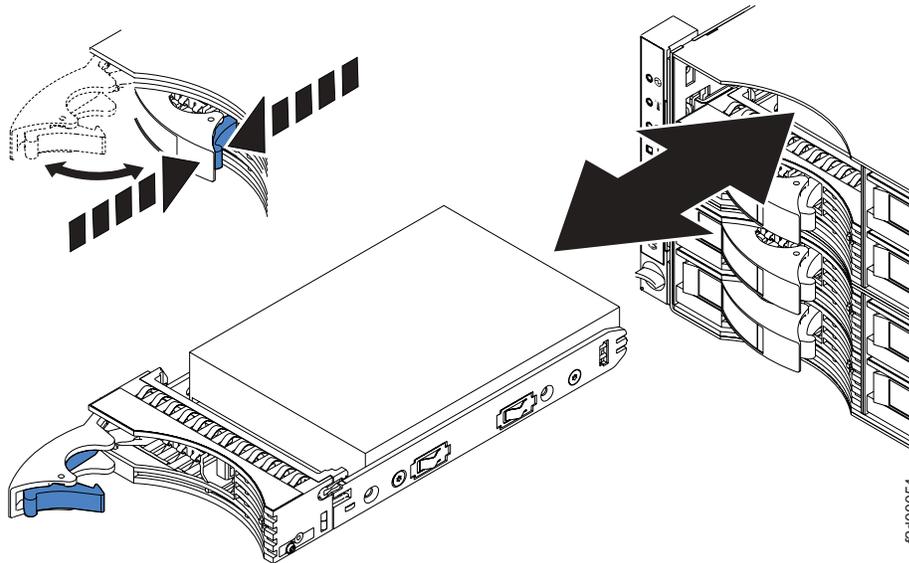


Note: If you are removing a defective disk drive module to replace it, be sure to handle the defective disk drive module carefully. Damage to defective disk drive modules can have a negative effect on failure analysis and warranty recovery.

To remove a disk drive module, perform the following steps:

1. Press the blue latch to release the disk drive module handle and pull the handle out and to your left. This action partially ejects the disk drive module out of its slot. Wait 30 seconds for the disk drive to spin down.
2. Grip the disk drive module with both hands to pull it toward you and completely out of the slot. Ensure that the disk drive is properly aligned in a horizontal position until the drive is clear of the enclosure. Failure to do so can result in physical damage to the drive or the drive component.

These images show a disk drive module removal procedure.



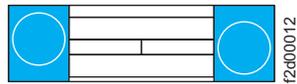
Important: All disk drive module slots must be filled to prevent overheating the storage unit. If necessary, insert a disk drive module blank in the empty slot.

Removing the power supply

Complete this task to remove the power supplies from the enclosure to reduce the weight of the enclosure during installation.

 Use approved ESD procedures to prevent damage.

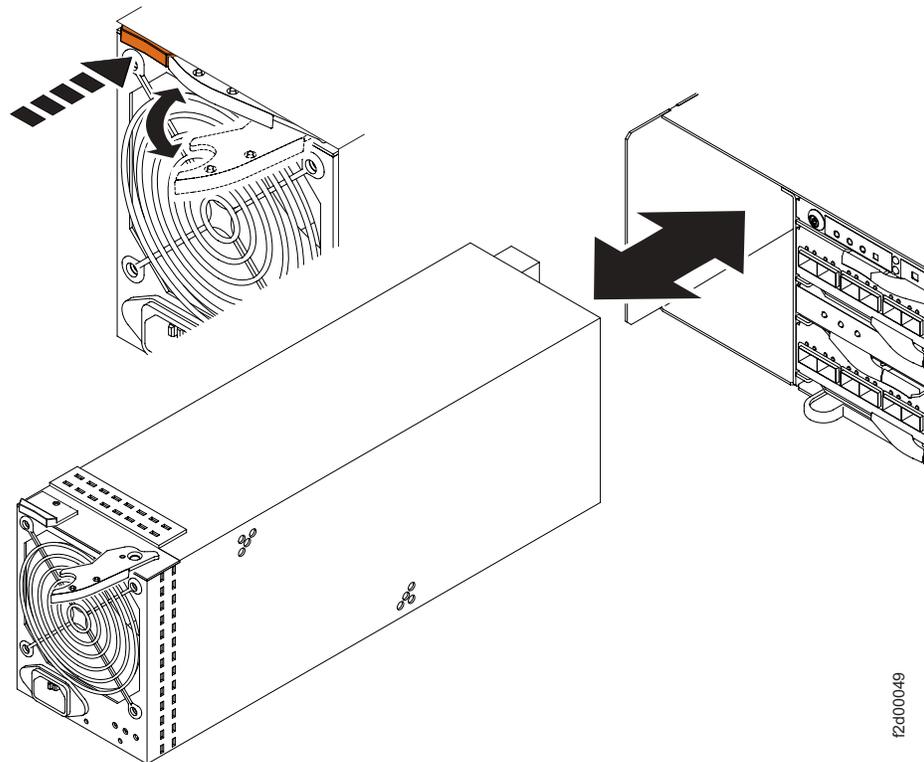
The power supplies are located on right and left sides of the rear of the enclosure.



To remove the power supplies, perform the following steps:

1. Press the release button. This releases the handle.
2. Pull the handle out and towards the middle of the enclosure. This action partially ejects the power supply from the slot.
3. Use the handle to pull the power supply partially out of the slot.
4. Grip the power supply with both hands to pull the unit completely from the slot.

These images show a power supply removal procedure.



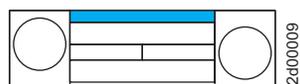
Removing the rear display panel

Complete this task to remove the rear panel display from the enclosure to reduce the weight of the enclosure during installation.



Use approved ESD procedures to prevent damage.

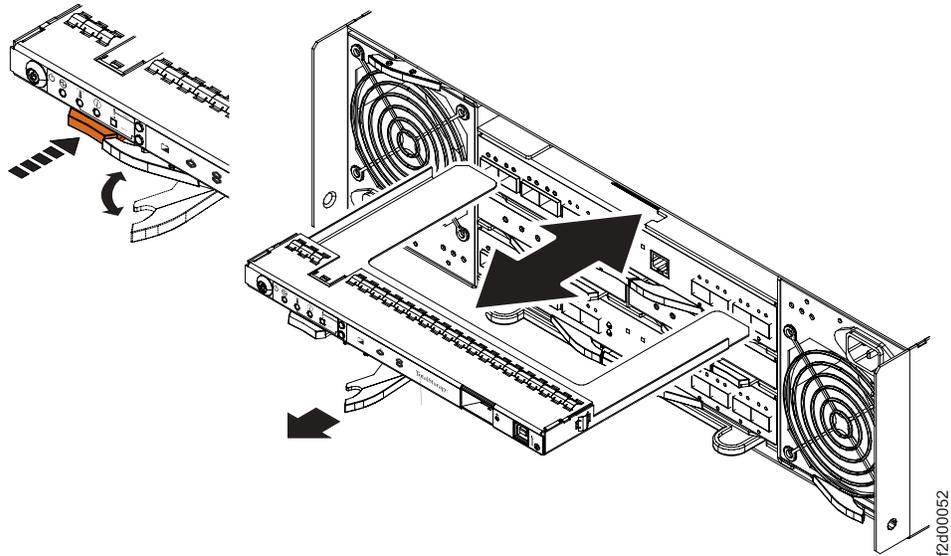
The rear operator panel is located at the rear of the expansion enclosure, on the top of the resource section.



To remove the rear panel display, perform the following steps:

1. Press the orange release button. This releases the handle.
2. Pull the handle out and to the right. This action partially ejects the display panel from the slot.
3. Use the handle to pull the rear display panel partially from the slot.
4. Grip the rear display panel with both hands to pull the unit completely from the slot.

These images show a rear display panel removal procedure.

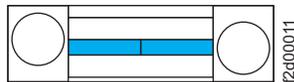


Removing the battery backup units

Complete this task to remove the battery backup units from the enclosure to reduce the weight of the enclosure during installation.

 Use approved ESD procedures to prevent damage.

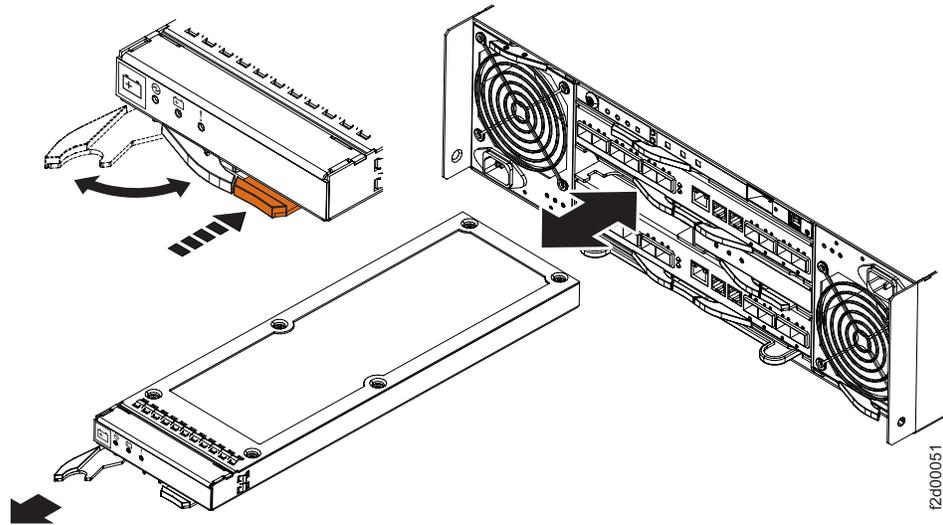
The battery backup units are located horizontally in the middle of the expansion enclosure.



To remove the battery backup units, perform the following steps:

1. Press the orange release button. This releases the handle.
2. Pull the handle out and to the left. This partially ejects the battery backup unit from the slot.
3. Use the handle to pull the battery backup unit partially out of the slot.
4. Grip the battery backup unit with both hands to pull it completely from the slot.

These images show a battery backup unit removal procedure.



Removing the processor cards

Complete this task to remove the processor cards from the enclosure to reduce the weight of the enclosure during installation.



Use approved ESD procedures to prevent damage.

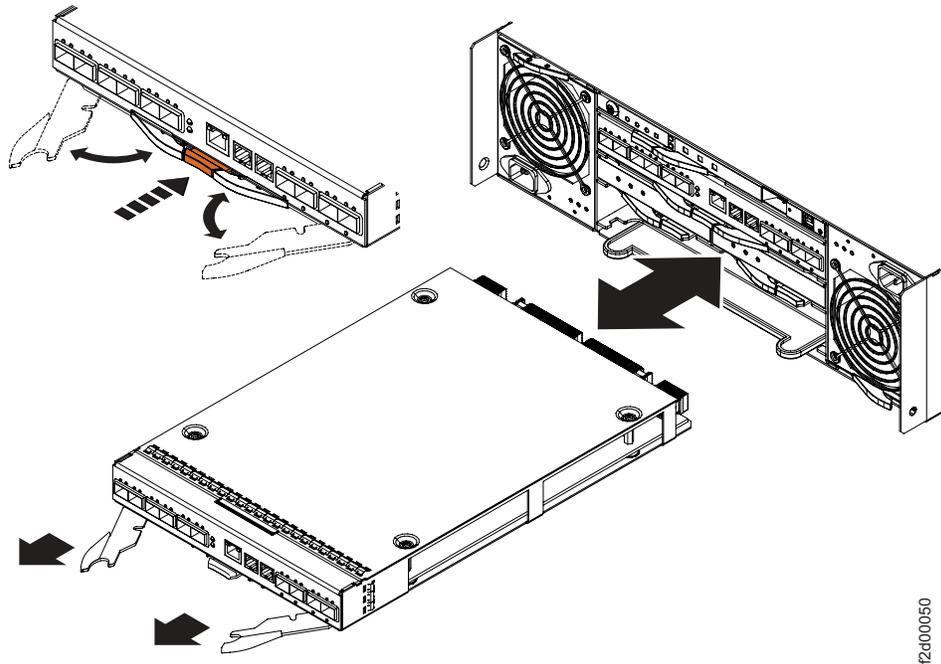
The processor cards are located in the rear of the expansion enclosure. One of the cards is located below the rear operator panel and above the battery backup units. The other card is located below the battery backup units and above the system service card.



To remove the processor cards, perform the following steps:

1. Press the orange release button. This releases both handles.
2. Grip both handles, one with each hand.
3. Pull both handles out and towards the outside of the enclosure. The right handle pulls out and to the right of the enclosure. The left handle pulls out and to the left of the enclosure. This partially ejects the processor card from the slot.
4. Use the handles to pull the processor card partially out of the slot.
5. Grip the processor card with both hands to pull the resource completely from the slot.

These images show a processor card removal procedure.



12d00050

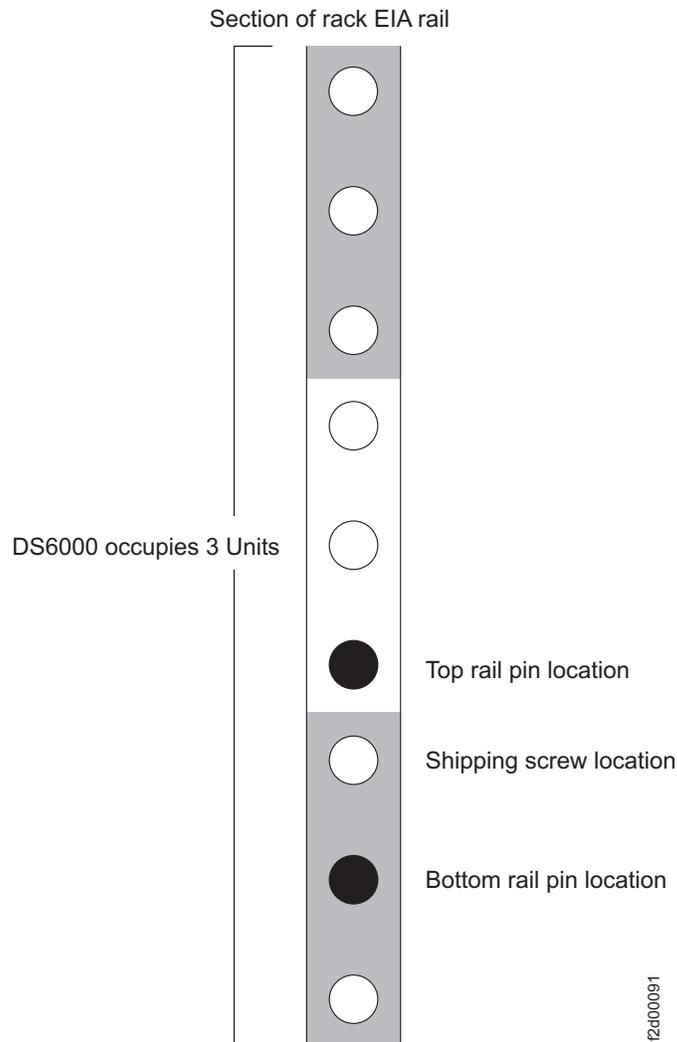
Installing the support rails

Complete this task to install the support rails in the rack.

The DS6000 requires one of the following supported IBM racks:

- 2101-200
- 7014
- 9308 (Netbay)
- Compatible Electronic Industries Association (EIA) 310-D Section 1 19-inch rack cabinet.

Use the following rack mounting template to identify the proper locations for inserting the slide rail pins.

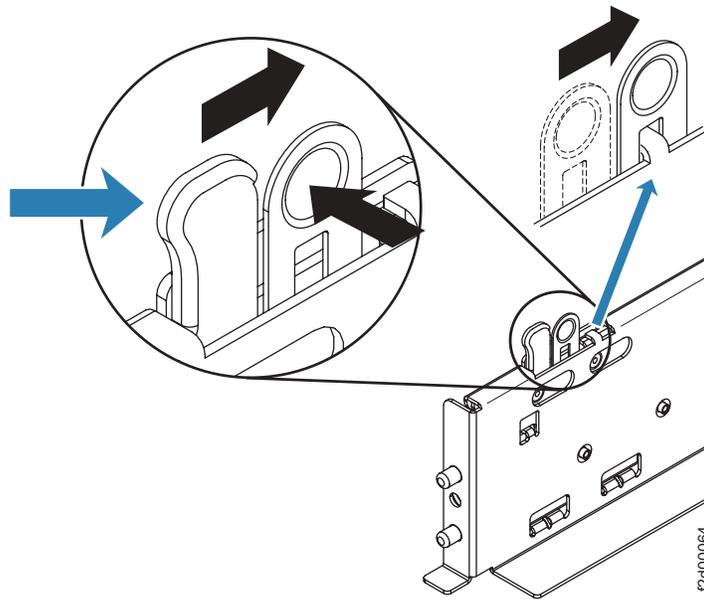


The distance between EIA rails, from the front to the rear of the rack, is 69.5 centimeters (27.36 inches) minimum to 76.5 centimeters (30.12 inches) maximum. This rack conforms to the EIA standard. Where you place the support rails in the rack depends on where you intend to position the server or expansion enclosure.

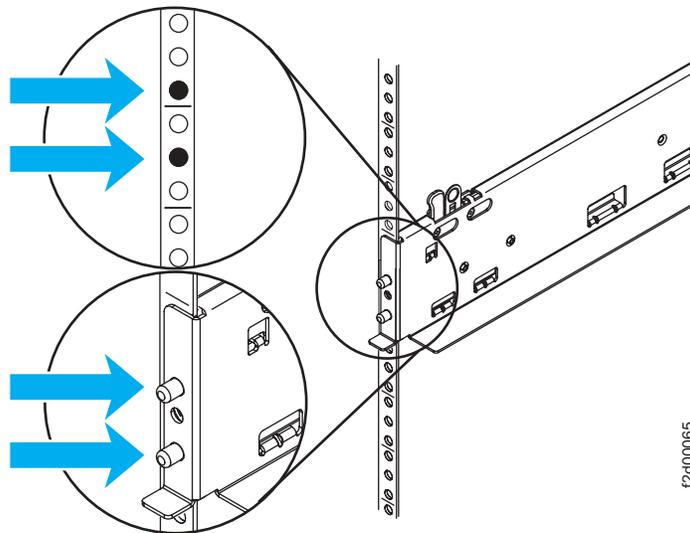
To install the support rails, perform the following steps:

Note: A set of rail pins and a mounting flange are on each end of the rail.

1. To set the rail pins in the open position, press firmly in on the slide-rail latch (right tab) while steadily pushing back on the finger pull (left tab) until the slide-rail latch is locked into position.



2. With rail pins open, fit the mounting flange around the rack cabinet rail. Align the rail pins with the holes in the rack cabinet rail and release the rail pins.



3. Repeat steps one and two with the other end of the rail, aligning with the rear mounting flange.
4. Repeat steps two and three with the left slide rail.

Installing the server enclosure in the rack

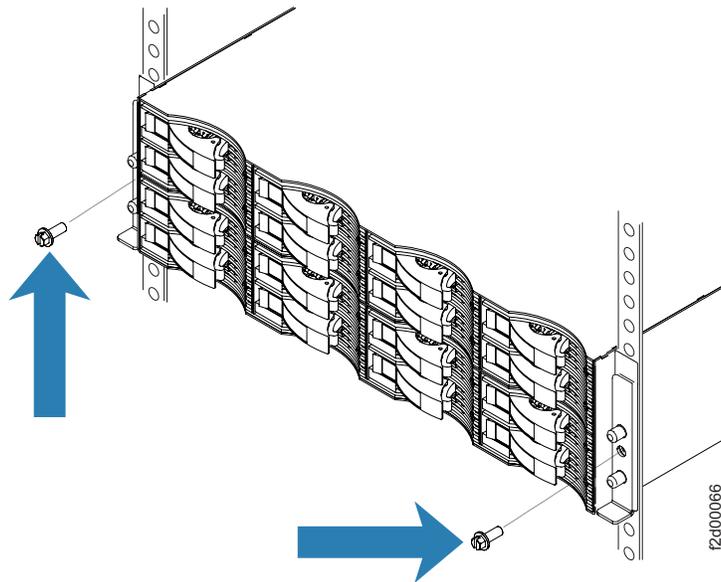
Complete this task to install the server enclosure in the rack.

Note: Screws are bagged and taped to the inside flange of the rail. Remove the bag prior to installing the server enclosures in the rack.

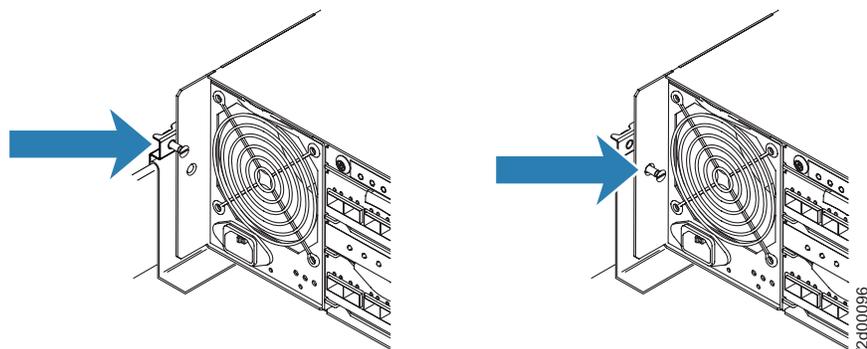
1. Place the server enclosure in the rack.

Note: If you removed the resources, you should be able to lift the unit into the cabinet with the help of one other person. If you did not remove the resources before installation, you need three people to lift the unit into the rack.

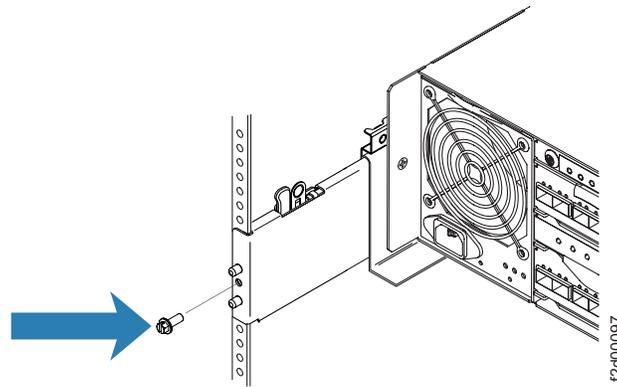
- a. Move the server enclosure to the front of the rack-mount cabinet.
- b. Slide the unit into the rack.
2. Secure the server enclosure to the rack.
 - a. Align the front mounting holes on each side of the server with the mounting holes on the front of the support rails.
 - b. Insert an M5 hex screw into the mounting holes between the two pins on both sides of the front of the unit. Tighten the screws to secure the front of the server enclosure to both of the front cabinet rails.



- c. Ensure the rear screw mounting bracket is pushed up against the rear of the chassis. Remove the M5 countersunk screw from the bracket and insert it into the hole in the rear of the chassis. Tighten the screws. Repeat for the opposite side.



- d. Insert an M5 hex screw into the mounting hole between the two pins on both sides of the rear of the chassis. Tighten the screws to secure the rear of the support rails to both of the rear cabinet rails.



Installing expansion enclosures in the rack

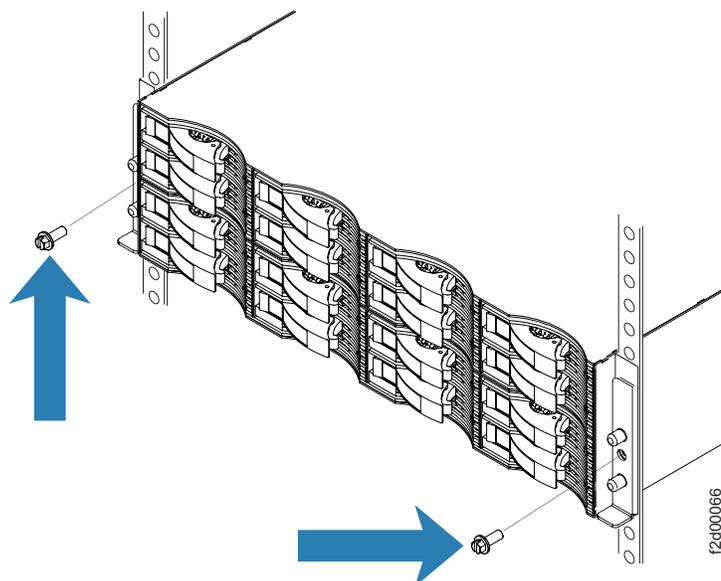
Complete this task to install the expansion enclosures in the rack.

Note: Screws are bagged and taped to the inside flange of the rail. Remove the bag prior to installing the expansion enclosures in the rack.

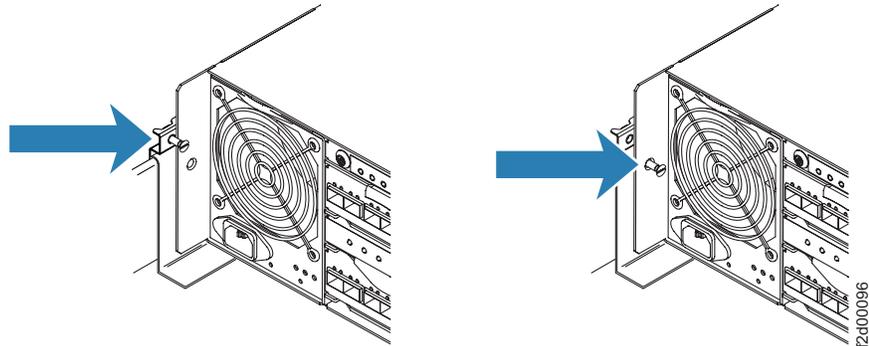
1. Place the expansion enclosure in the rack.

Note: If you removed the resources, you should be able to lift the unit into the cabinet. If you did not remove the resources before installation, you should have three people that are available to lift the unit into the rack.

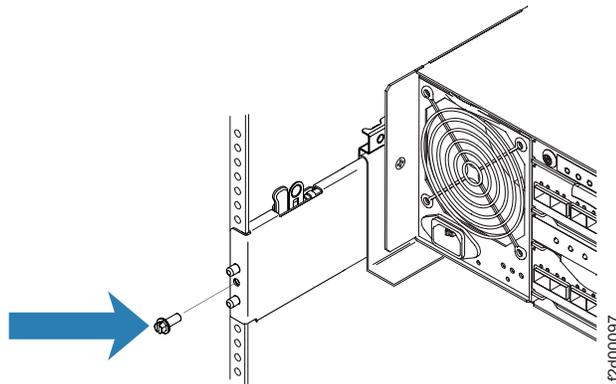
- a. Move the expansion enclosure to the front of the rack-mount cabinet.
 - b. Slide the unit into the rack.
2. Secure the expansion enclosure to the rack.
 - a. Align the front mounting holes on each side of the server with the mounting holes on the front of the support rails.
 - b. Insert an M5 hex screw into the mounting hole between the two pins on both sides of the front of the unit. Tighten the screws to secure the front of the expansion enclosure to both of the front cabinet rails.



- c. Ensure the rear screw mounting bracket is pushed up against the rear of the chassis. Remove the M5 countersunk screw from the bracket and insert it into the hole in the rear of the chassis. Tighten the screw. Repeat for the opposite side.



- d. Insert an M5 hex screw into the mounting holes between the two pins on both sides of the rear of the chassis. Tighten the screws to secure the rear of the support rails to both of the rear cabinet rails.



Replacing the hardware resources

This section contains information that helps you replace the hardware resources in the enclosure after installing the enclosure in the rack. If you did not remove the hardware resources before installation, you can skip this section.

1. Make sure to follow ESD procedures by using an anti-static wrist strap.
2. Replace all the hardware resources in the DS6000 expansion and server enclosures that were previously removed.

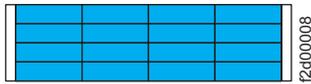
Replacing the disk drive module

If you removed the disk drive modules from the enclosure to reduce the weight of the enclosure before installation, complete this task to replace the disk drive modules in the enclosure.



Use approved ESD procedures to prevent damage.

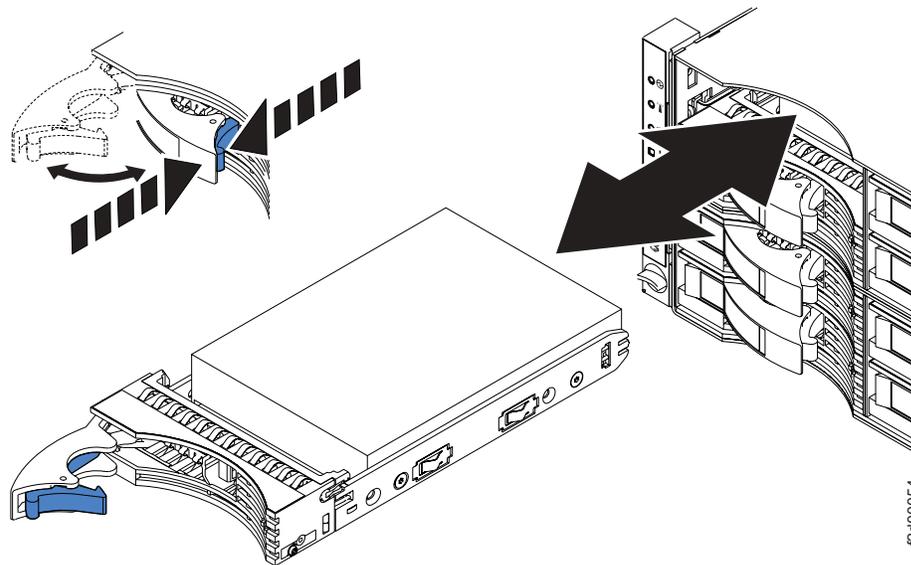
The disk drive modules are located in rows and columns on the front of the storage unit.



To replace the disk drive modules, perform the following steps:

1. Before installing the disk drive module, open the disk drive module handle by pressing the blue latch and pulling the handle open.
2. Align the disk drive module with the groove on the disk drive module bay and push it into its slot. The drive stops before it is fully seated. Ensure that the disk drive is properly aligned in a horizontal position. Failure to do so could result in physical damage to the drive or the drive component.
3. Push the disk drive module handle to the right until it is latched closed.
4. Verify that the front of the disk drive module is aligned with the other disk drive modules.

These images show a disk drive module replacement procedure.



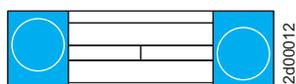
Important: All disk drive module slots must be filled to prevent overheating the storage unit. If necessary, insert a disk drive module blank in an empty slot.

Replacing the power supply

If you removed the power supplies to reduce the weight of the enclosure before installation, complete this task to replace the power supplies.

 Use approved ESD procedures to prevent damage.

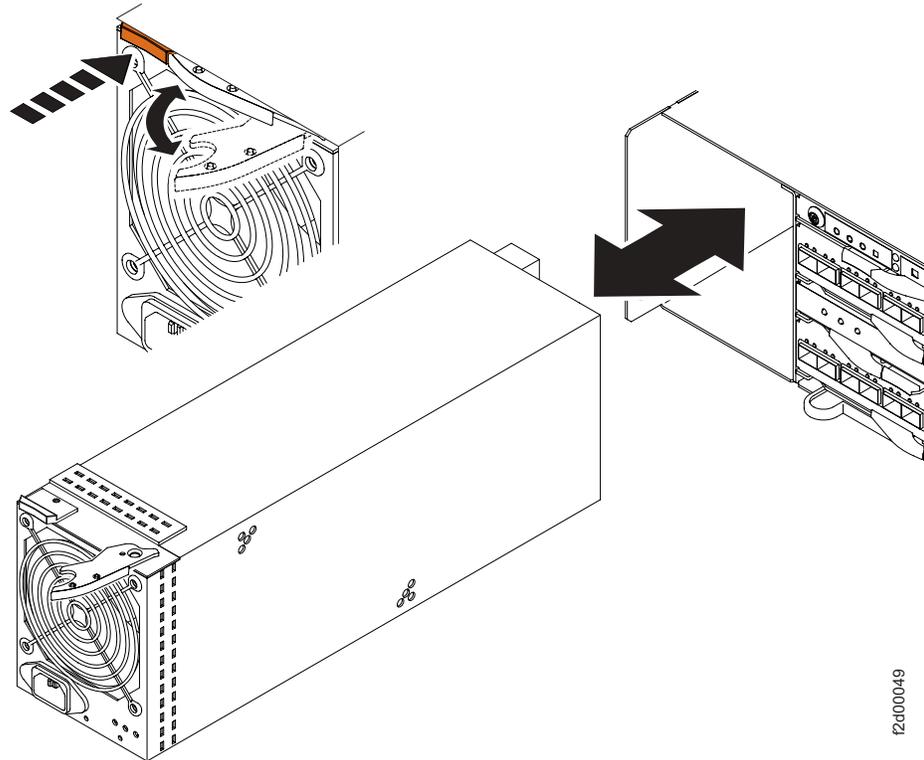
The power supplies are located on the right and left sides of the rear of the enclosure.



1. Grip the power supply with both hands and align the resource with the slot.

2. Using both hands, push the power supply into the slot until the handle mechanism stops the forward movement.
3. Push the handle in and toward the outer edge of the enclosure until the release button clicks. This inserts the remaining portion of the power supply into the slot.

These images show a power supply replacement procedure.



f2d00049

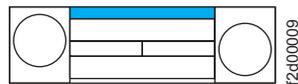
Replacing the rear display panel

If you removed the rear display panel to reduce the weight of the enclosure before installation, complete this task to replace the rear display panel.



Use approved ESD procedures to prevent damage.

The rear display panel is located at the top of the rear of the enclosure.

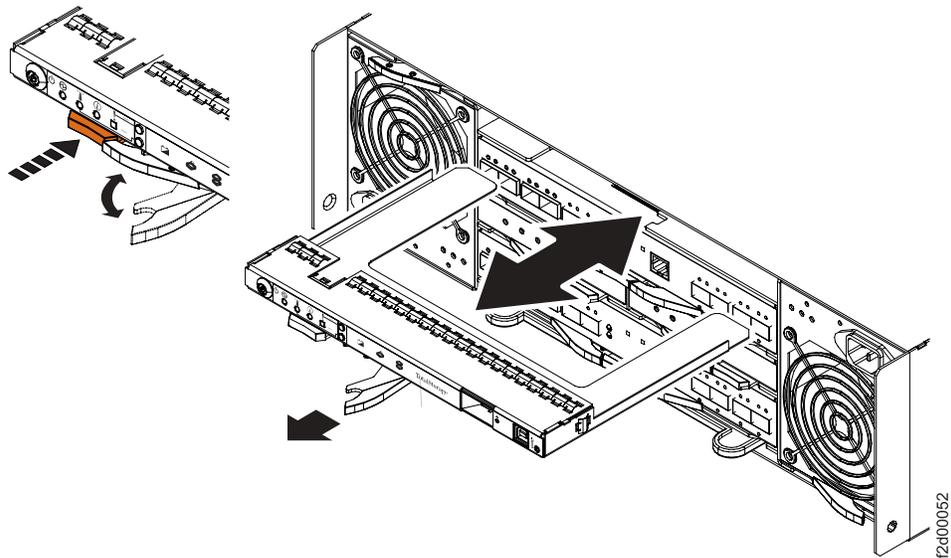


f2d00009

To replace the rear display panel, perform the following steps:

1. Grip the rear operator panel with both hands and align the resource with the slot.
2. Using both hands, push the rear operator panel into the slot until the handle mechanism stops the forward movement.
3. Push the handle in and to the left until the release button clicks. This inserts the remaining portion of the rear operator panel into the slot.

These images show a rear operator panel replacement procedure.

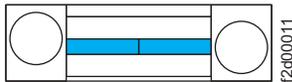


Replacing the battery backup unit

If you removed the battery backup units to reduce the weight of the enclosure before installation, complete this task to replace the battery backup units.

 Use approved ESD procedures to prevent damage.

The battery backup units are located horizontally in the middle of the rear of the enclosure.

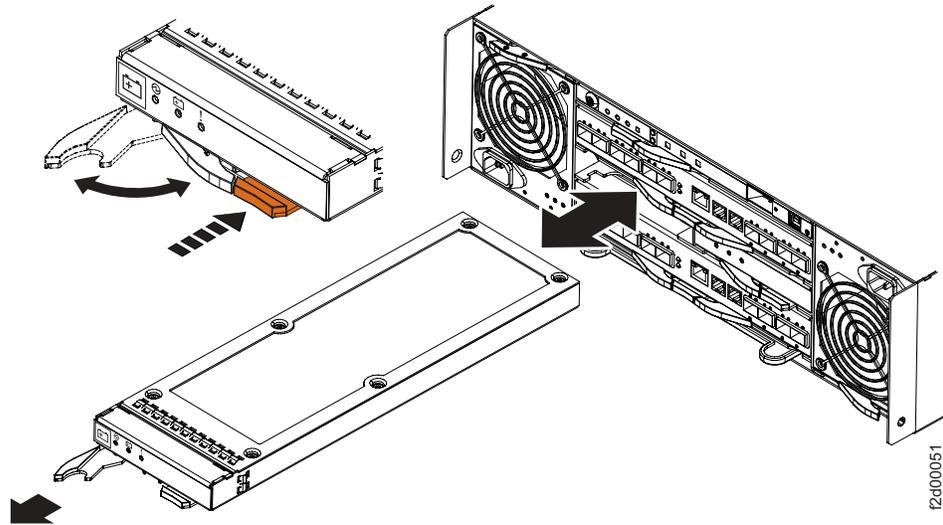


To replace the battery backup unit, perform the following steps:

1. Grip the battery backup unit with both hands and align the resource with the slot.
2. Using both hands, push the battery backup unit into the slot until the handle mechanism stops the forward movement.
3. Push the handle in and to the right until the release button clicks. This inserts the remaining portion of the battery backup unit into the slot.

The battery backup unit automatically resumes its function after the batteries are fully charged.

These images show a battery backup unit replacement procedure.



Replacing the processor cards

If you removed the processor cards to reduce the weight of the enclosure before installation, complete this task to replace the processor cards.



Use approved ESD procedures to prevent damage.

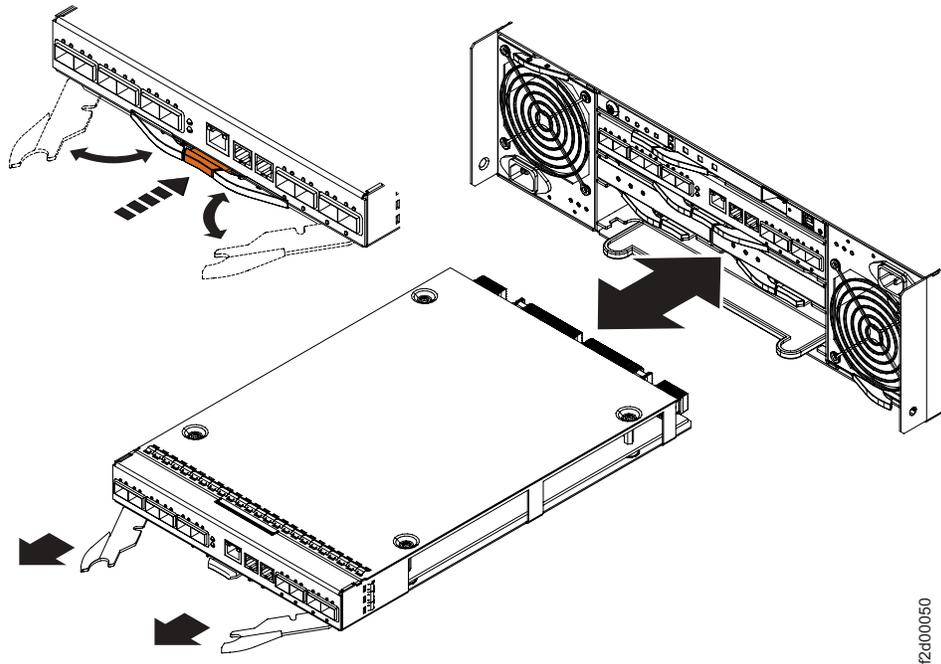
The processor cards are located in the rear of the enclosure. One of the cards is below the rear operator panel and above the battery backup units. The other card is located below the battery backup units and above the system service card.



To replace the processor cards, perform the following steps:

1. Grip the processor card with both hands to align the resource with the slot.
2. Using both hands, push the processor card into the slot until the handle mechanism stops the forward movement.
3. Push both handles in and towards the center of the enclosure at the same time until the release button clicks. This inserts the remaining portion of the processor card into the slot.

These images show a processor card replacement procedure.



12d00050

Installing the front display panel

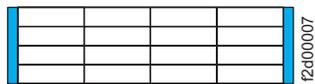
Complete this task to install or replace the front display panel or the right front bezel. If you are installing the front panel display after removing it for a servicing procedure, you can install it whether the system is powered off or powered on.



Use approved ESD procedures to prevent damage.

Important: When you replace the front display panel in an enclosure, ensure that all resources from a particular server or expansion enclosure are reinstalled into the same physical enclosure and in the same location inside that enclosure.

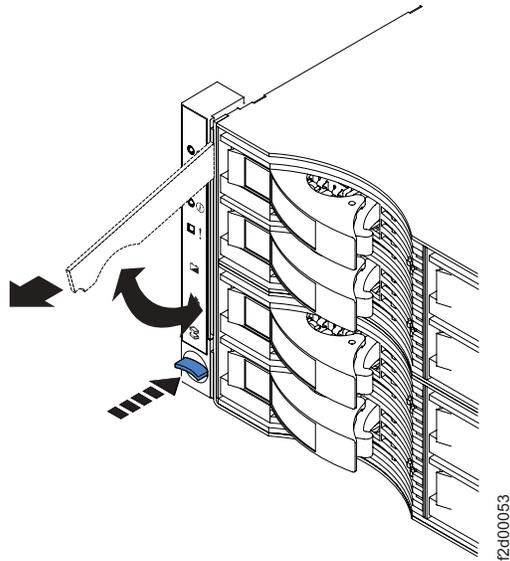
The front display panel and right front bezel are located on the left and right of the front of the enclosure, respectively.



To install the front panel display, perform the following steps:

1. Use the handle to align the top of the display panel with the empty slot. The top of the front display panel must fit over the metal support at the top of the slot at the same time that the back of the handle fits under the metal catch on the side of the slot.
2. Push the handle down and into the recessed slot so that it is flush with the front of the enclosure. This will lock the front display panel into position.

This image shows a front display panel replacement procedure.



Installing the host systems and I/O adapters

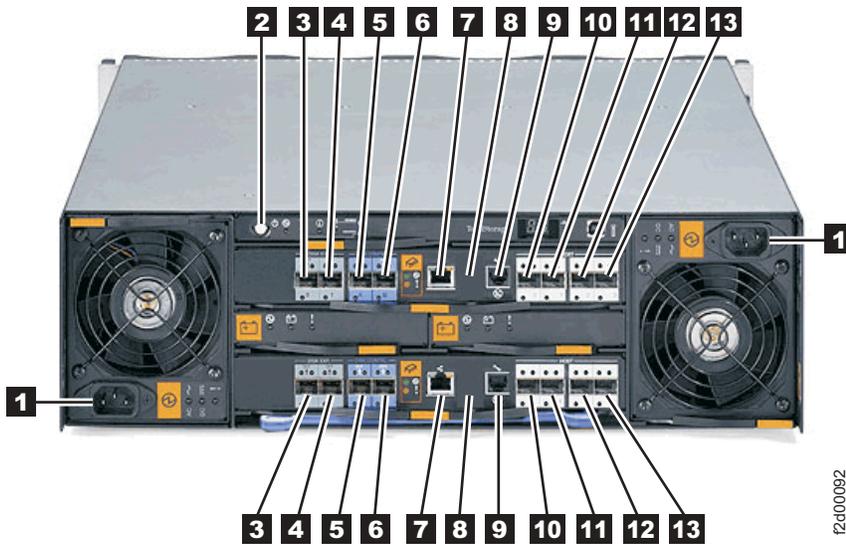
Complete this task to install the host systems and I/O adapters.

Note: See the *IBM System Storage DS6000 Host Systems Attachment Guide* and the documentation provided with your I/O adapters for installation requirements and procedures.

1. Use the correct I/O adapter driver. For the latest supported I/O adapters and drivers, go to the interoperability matrix at the following Web site: <http://www.ibm.com/servers/storage/disk/ds6000/interop.html>.
2. A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www-03.ibm.com/servers/storage/support/config/hba/index.wss>
3. Attach fiber optic interface cables to each I/O adapter. You will connect the other end of the cables to the server enclosure later in the installation process.

Routing the cables

This section provides information on routing cables. The following figure shows an overview of all of the connections in the server enclosure.



1. Power supply connector
2. Power button
3. Server enclosure processor card Disk Exp Port 0
4. Server enclosure processor card Disk Exp Port 1
5. Server enclosure processor card Disk Ctlr Port 0
6. Server enclosure processor card Disk Ctlr Port 1
7. Server enclosure processor card Ethernet port
8. Server enclosure processor card SCSI Enclosure Services (SES) serial port
9. Server enclosure processor card symmetric multiprocessing (SMP) serial port
10. Server enclosure processor card Host Port 0
11. Server enclosure processor card Host Port 1
12. Server enclosure processor card Host Port 2
13. Server enclosure processor card Host Port 3

DS6000 fabric zoning

Creating zones for your host attachment port connections enables you to manage traffic and establish redundancy that eliminates single points of hardware failure.

Depending on your environment, you can benefit from isolating traffic as much as possible in your storage area network. Networks with a large amount of volumes and heavy host traffic can benefit the most. Implementing single-initiator, single-target zoning allows you to isolate traffic for each port. Single-initiator, single-target zoning creates small zones within the fabric with only two zone members (ports). The zone consists of one target (a storage unit port), and one initiator (a host system port). The key benefit of single-initiator, single-target zoning is traffic isolation or masking.

Figure 2 on page 55 displays two host systems that are attached to a DS6800 through two fabric switches. Through the switches, each of the host systems is connected to both of the DS6800 processor cards (PC-0 and PC-1) using two host adapters in each host system (HA-1 and HA-2). Traffic isolation and masking occur by creating a zone for each host system port.

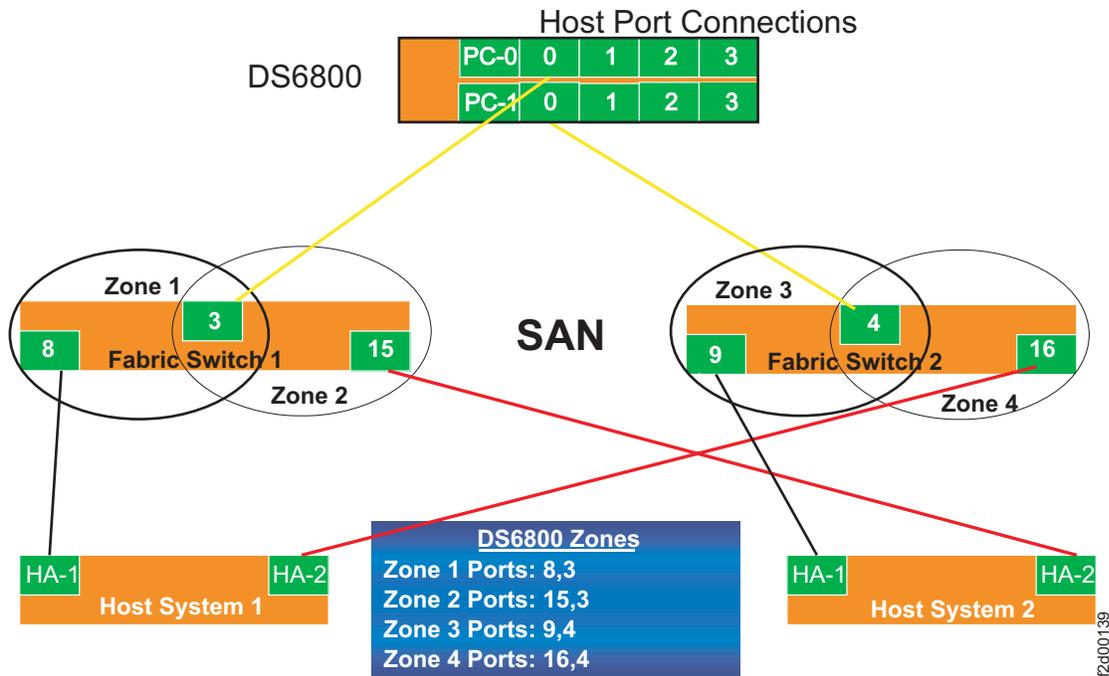


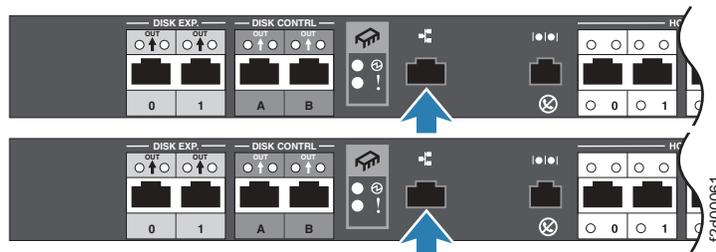
Figure 2. DS6000 zoning example

Connecting network cables

Complete this task to either connect your server enclosure to your local area network or connect your processor cards to each other. This enables communication between the processor cards that is essential for correct operation.

To connect the network cables, perform the following steps:

1. Use the Ethernet interface port on the back of the enclosure to connect the processor cards to your external Ethernet network (Ethernet switch) for direct management of the server enclosure.



2. The serial interface ports are intended to be used by service technicians to perform diagnostic operations on the server enclosure. Incorrect use of the serial port can result in loss of data access and, in some cases, in loss of data. The processor cards must be able to communicate with each other when the server enclosure is powered on. If you do not connect the server enclosure to the local area network, you must enable communication between the two processor cards and the management console. Use an Ethernet crossover cable to connect the two processor cards to each other and the serial cable that was shipped with the enclosure to connect one of the processor cards to the management console. You can use this temporary connection to set the IP address on each of the processor cards.

Installing the fiber optic cables and SFPs

Complete this task to complete the initial installation of fiber optic cables and small form factor pluggables (SFPs).



Use approved ESD procedures to prevent damage.

Before you install the fiber optic cables, you must install either a shortwave or longwave SFP in each port that you use to create an interenclosure or host connection.

The fiber optic ports are located on the processor card in the enclosure.



1. Lift the wire or plastic tab and remove the protective cap from the SFP module.
2. Remove the protective cap from the SFP module port on the server enclosure.
3. Insert the SFP module into the SFP module port on the server enclosure.
4. Remove the protective caps from the fiber optic cable.
5. Insert the fiber optic cable into the SFP module.
6. Check to ensure that both the SFP module and fiber optic cable are properly seated in the port.

Connecting expansion enclosures

Complete this task to connect expansion enclosures for the DS6000.

Use the following procedures to connect your expansion enclosures during a new installation. When expansion enclosures are connected during a new installation, these procedures are nonconcurrent (the DS6000 is turned off). If you are adding an expansion enclosure to an existing storage unit, see "Adding an expansion enclosure" in the Information Center at <http://publib.boulder.ibm.com/infocenter/ds6000ic/index.jsp>.

Note: You can add multiple expansion enclosures at one time. However, you must use the following steps to maintain proper loop balancing.

The recommended sequence for adding up to seven expansion enclosures is as follows:

1. Connect the first and second enclosures on Loop 1.
2. Connect the third enclosure on Loop 0.
3. Alternate each additional expansion enclosure between Loop 1 and Loop 0, starting with Loop 1.

For example, the first two expansion enclosures are on Loop 1, the third expansion enclosure is on Loop 0, and so on. You can connect up to four expansion enclosures on Loop 1 and up to three on Loop 0, for a total of seven expansion enclosures.

Expansion enclosure	Loop
1	1
2	1

Expansion enclosure	Loop
3	0
4	1
5	0
6	1
7	0

Figure 3 provides a logical view of the loop configuration for attached expansion enclosures.

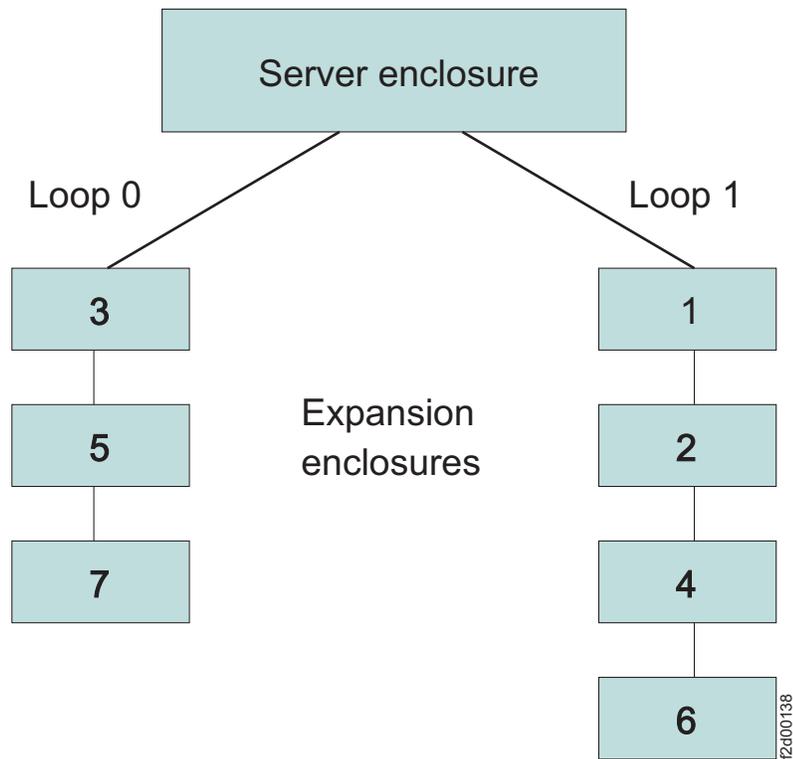


Figure 3. Connecting expansion enclosures

Figure 4 on page 58 shows a recommended rack configuration for a server enclosure, management console (MC), and seven expansion enclosures (labeled EXP) on Loop 1 and Loop 0. The server enclosure is located in the middle of the rack, with the management console immediately above it. The expansion enclosures for Loop 1 are installed above the management console, in descending order from the top down, with expansion enclosure 6 at the top of the rack and expansion enclosure 1 immediately above the management console. The expansion enclosures for Loop 0 are installed below the server enclosure, in ascending order, with expansion enclosure 3 immediately below the server enclosure and expansion enclosure 7 at the bottom of the rack.

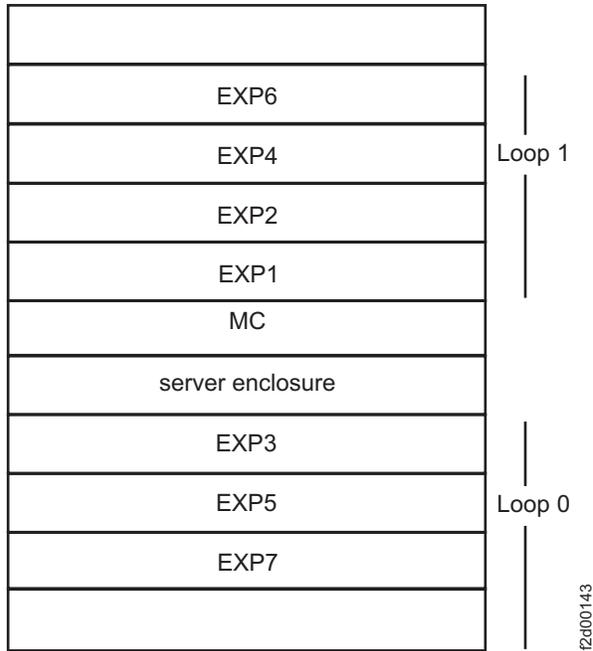


Figure 4. Recommended rack configuration

Attention: Ensure that you have planned the layout and have installed all your expansion enclosures before you start connecting them to their appropriate loops. Because Loop 1 and Loop 0 have different physical cabling connections, you must follow each step exactly.

1. Connect the expansion enclosures to a server enclosure on Loop 1, using the following tables.

The following tables provide a step-by-step outline that you can use to properly connect your expansion enclosures to the server enclosure. Table 5 provides an outline of the initial connection between the server enclosure and expansion enclosure 1 on Loop 1. Table 6 on page 59 outlines a single connection between two expansion enclosures on Loop 1, such as the connection that occurs between expansion enclosures 1 and 2, expansion enclosures 2 and 4, or expansion enclosures 4 and 6. You must complete the connection in Table 6 on page 59 for each expansion enclosure after expansion enclosure 1. Each row in a table requires a single fiber-optic cable and represents a single physical connection between two ports.

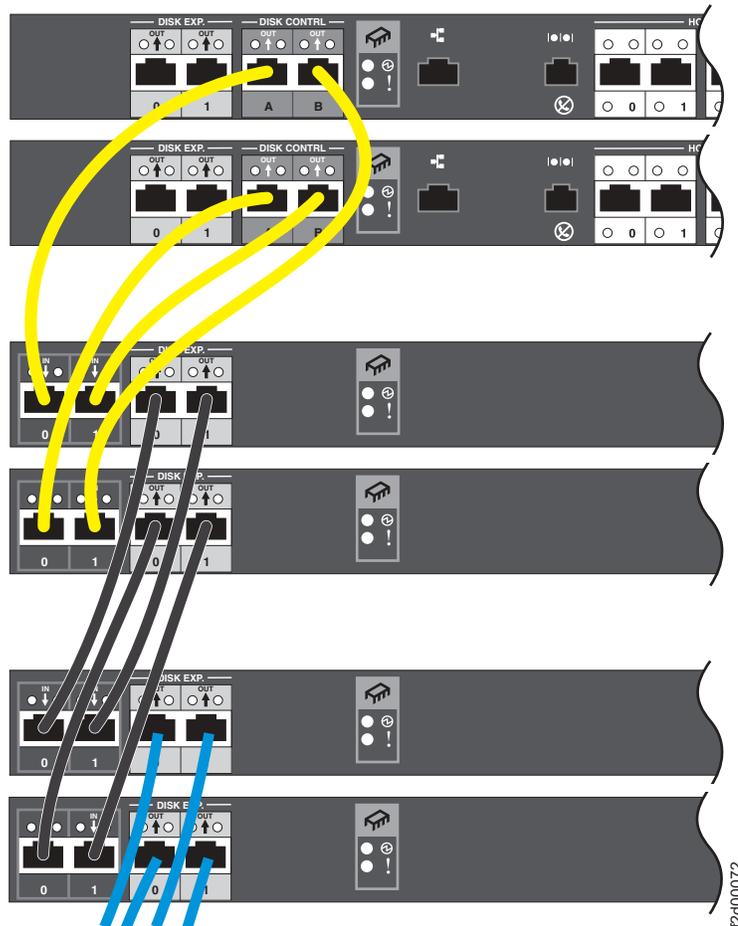
Table 5. Loop 1 initial connection between the server enclosure and expansion enclosure 1

From	On	To	On
left Disk Contrl OUT port	upper server enclosure processor card	left IN port	upper expansion enclosure processor card
right Disk Contrl OUT port	upper server enclosure processor card	right IN port	lower expansion enclosure processor card
left Disk Contrl OUT port	lower server enclosure processor card	left IN port	lower expansion enclosure processor card
right Disk Contrl OUT port	lower server enclosure processor card	right IN port	upper expansion enclosure processor card

Table 6. Loop 1 connection between expansion enclosure 1 and expansion enclosure 2 and subsequent connections between expansion enclosures 2, 4, and 6

From	On	To	On
left OUT port	upper expansion enclosure processor card of the first expansion enclosure	left IN port	upper expansion enclosure processor card of the second expansion enclosure
right OUT port	upper expansion enclosure processor card of the first expansion enclosure	right IN port	upper expansion enclosure processor card of the second expansion enclosure
left OUT port	lower expansion enclosure processor card of the first expansion enclosure	left IN port	lower expansion enclosure processor card of the second expansion enclosure
right OUT port	lower expansion enclosure processor card of the first expansion enclosure	right IN port	lower expansion enclosure processor card of the second expansion enclosure

Attention: Expansion enclosures that are not cabled correctly can cause problems at initialization and can cause unexpected error log entries. Check the following figure carefully when you are connecting the cables.



2. Connect the expansion enclosures to a server enclosure on Loop 0, using the following tables.

The following tables provide a step-by-step outline that you can use to properly connect your expansion enclosures to the server enclosure. Table 7 provides an outline of the initial connection between the server enclosure and expansion enclosure 3 on Loop 0. Table 8 outlines a single connection between two expansion enclosures on Loop 0, such as the connection that occurs between expansion enclosures 3 and 5 or expansion enclosures 5 and 7. You must complete the connection in Table 8 for each expansion enclosure after expansion enclosure 3. Each row in a table requires a single fiber-optic cable and represents a single physical connection between two ports.

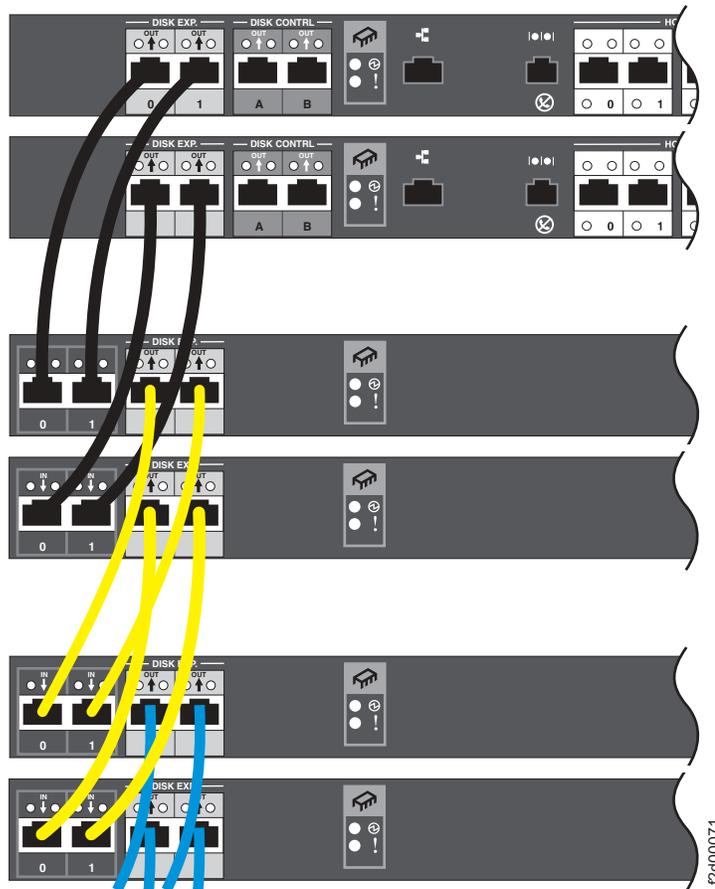
Table 7. Loop 0 initial connection between the server enclosure and expansion enclosure 3

From	On	To	On
left Disk Exp OUT port	upper server enclosure processor card	left IN port	upper expansion enclosure processor card
right Disk Exp OUT port	upper server enclosure processor card	right IN port	upper expansion enclosure processor card
left Disk Exp OUT port	lower server enclosure processor card	left IN port	lower expansion enclosure processor card
right Disk Exp OUT port	lower server enclosure processor card	right IN port	lower expansion enclosure processor card

Table 8. Loop 0 connection between expansion enclosure 3 and expansion enclosure 5 and subsequent connections between expansion enclosures 5 and 7

From	On	To	On
left OUT port	upper expansion enclosure processor card of the first expansion enclosure	left IN port	upper expansion enclosure processor card of the second expansion enclosure
right OUT port	upper expansion enclosure processor card of the first expansion enclosure	right IN port	upper expansion enclosure processor card of the second expansion enclosure
left OUT port	lower expansion enclosure processor card of the first expansion enclosure	left IN port	lower expansion enclosure processor card of the second expansion enclosure
right OUT port	lower expansion enclosure processor card of the first expansion enclosure	right IN port	lower expansion enclosure processor card of the second expansion enclosure

Attention: Expansion enclosures that are not cabled correctly can cause problems at initialization and cause unexpected error log entries. Check the following figure carefully when you are connecting the cables.



Connecting hosts to server enclosure processor cards

Complete this task to connect the host adapter to the server enclosure processor cards.

1. Install an SFP in a host port on the server enclosure processor card.
2. Connect the host system cables to the server enclosure processor card or to switches. Figure 5 shows the location on the processor cards where the host system cables connect.



Figure 5. Connecting host system cables to processor cards

3. Repeat steps 1 and 2 as necessary.

Use the examples in the following diagrams to connect hosts to the DS6000 using switches. When you use external fibre-channel switches, the DS6000 supports up to 1040 host attachments.

Figure 6 shows an example of a dual path configuration using fibre-channel switches. Host 1 contains two host bus adapters that are connected to the server enclosure processor cards. To configure a host with dual path redundancy, connect the first host bus adapter (HA1) to SW1 and connect the second host bus adapter (HA2) to SW2. Then, connect SW1 to the upper server enclosure processor card and SW2 to the lower server enclosure processor card.

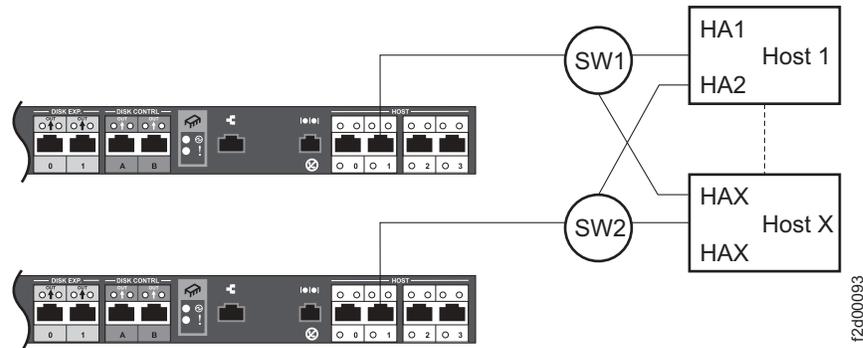


Figure 6. Dual fibre-channel switched path configuration

You can directly attach up to four fibre-channel switches to a DS6000. Figure 7 shows an example of attaching four fibre-channel switches.

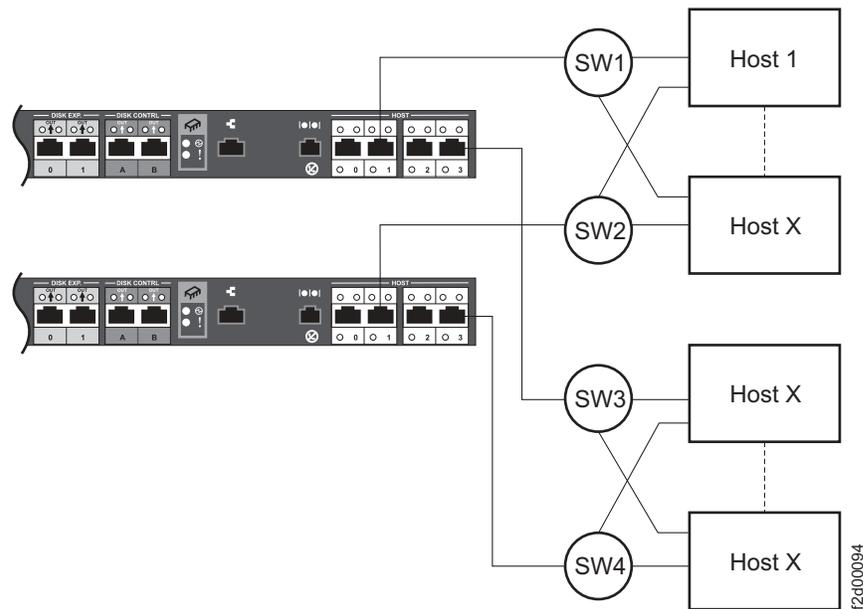


Figure 7. Dual path configuration with four fibre-channel switches

If the configuration requires more hosts than four fibre-channel switches can support, you must add cascading switches, as shown in Figure 8 on page 63.

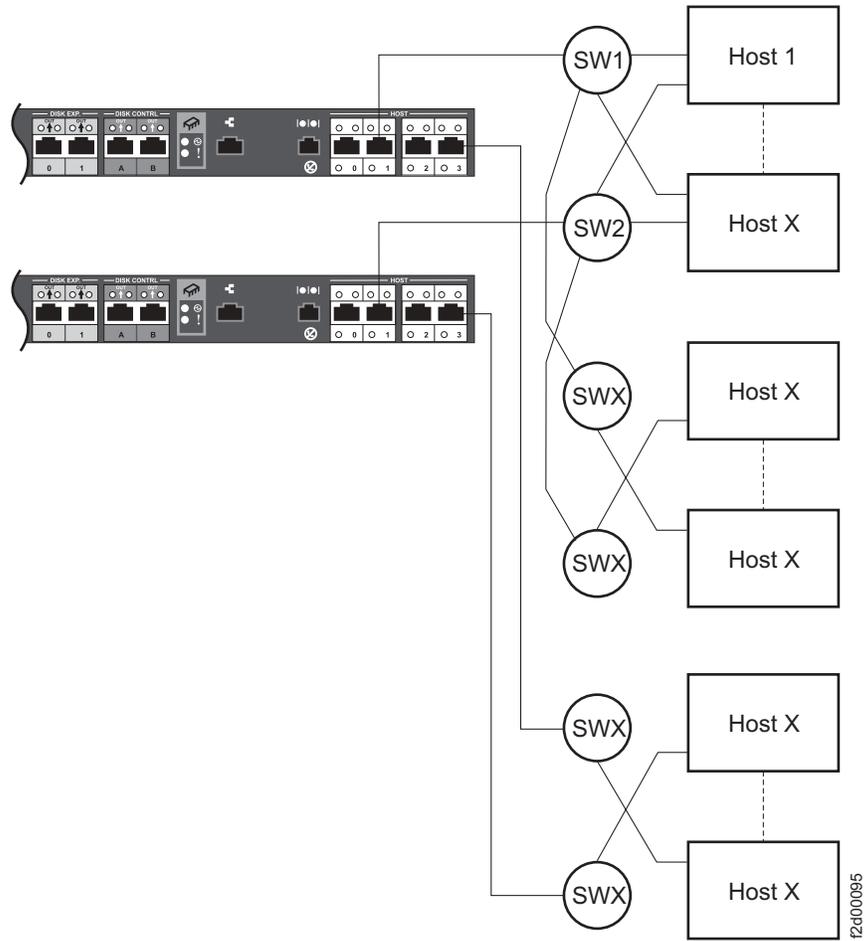


Figure 8. Cascading fibre-channel switched configuration

Connecting power cords

Complete this task to connect the power cords. You can connect the power cords to a primary power unit inside the rack, such as a properly grounded ac distribution unit, or to an external source, such as a properly grounded electrical outlet.

Each server or expansion enclosure uses two standard power cords.

To connect the power cords, perform the following steps:

1. If you have not already done so, attach the power cords to the server and expansion enclosures. Figure 9 on page 64 shows the power cord locations.

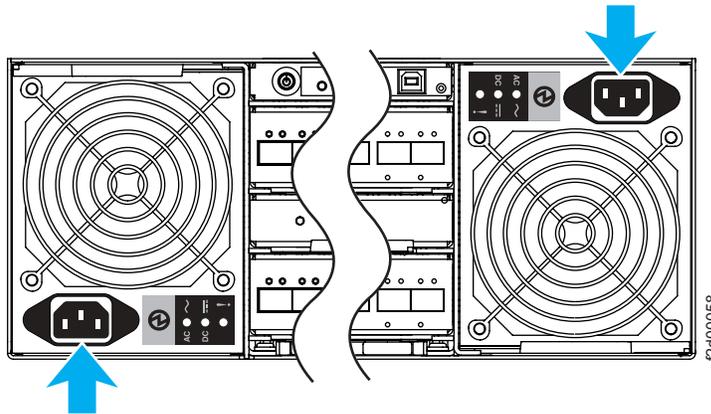


Figure 9. Power cord locations for the server and expansion enclosures

2. Connect the other end of the right server enclosure power cord to a power supply, by plugging it into a properly grounded electrical outlet.
3. Connect the other end of the left server enclosure power cord to a power supply, by plugging it into a properly grounded electrical outlet. To maintain power redundancy, plug it into a separate independent external power circuit.
4. If you have expansion enclosures, repeat the above steps for each expansion enclosure. Remember to plug the right and left power supplies expansion enclosures into the same two independent external power circuits.

Powering on

Complete this task to power on a storage unit.

Before powering on a storage unit, ensure that the management console has been installed. See Chapter 4, “Installing the DS6000 Storage Manager,” on page 73.

Perform the following steps to turn on the power for the initial startup of the storage unit:

1. Verify that all communication, interenclosure connections, and power cables are plugged into the back of the enclosures and into properly grounded ac electrical outlets.
2. Verify that all disk drive modules are locked securely in place.

Note: Ensure that at least four disk drive modules are present before you start the storage unit.

3. Check the system documentation for the external hardware devices that you intend to power on, and then determine the proper startup sequence. Use the following power-on sequence, where it is applicable:
 - a. Turn on the power to the supporting devices (for example, Ethernet switches and management stations) before turning on the server enclosure.
 - b. Turn on the server enclosure by pressing the power button that is located on the left side of the rear operator panel.

Notes:

- 1) The unit powers on the server enclosure and then automatically powers on the attached expansion enclosures in the necessary sequence. After the server enclosure has completely powered on, the power-on sequence continues down the line of attached

expansion enclosures. Each expansion enclosure powers on the enclosed resources in the necessary order before the power-on sequence continues to the next attached expansion enclosure. Enclosures on loop 1 might power on after the enclosures that are connected to loop 0, and they can take several minutes to power on.

- 2) If you are powering on a new enclosure or if the enclosure has been powered down long enough for the battery backup unit (BBU) to lose its charge, the BBU must build at least a 50% charge before the storage unit is fully functional. This can take up to two and one-half hours. The storage unit might not be able to communicate with the DS Storage Manager and host systems until the BBUs have accumulated a sufficient charge.
- 3) After powering on the enclosure, the information indicator flashes to indicate that the system is initializing and is not yet ready to be accessed by either the management console or host systems. Once the flashing amber light turns off, initialization is complete and the system is ready for I/O.

CAUTION:

Upon power on, a discovery process is initiated to detect all hardware. This process can take an extensive amount of time, depending on the number and type of attachments. You must wait until this process is complete before initiating any other processes.

Verifying status through LEDs

The LEDs display the status of the expansion server and components. Green LEDs indicate a normal operating status; amber LEDs indicate a possible failure.

Note: The green drive active LED and amber drive fault LED might flash intermittently as the drives spin-up. Wait until the DS6000 is finished powering up before you check the LEDs on the front of the storage unit.

After you power on the enclosure, the information indicator flashes to indicate that the system is initializing and is not yet ready to be accessed by either the management console or host systems. Initialization is complete, and the system is ready for I/O, once the flashing amber light turns off.

To verify the status of the server and expansion enclosures, perform the following steps:

1. At the front and rear of the server and expansion enclosures, verify that the LEDs are displaying the appropriate state. Table 9 provides a list of LED information. Additional LED information is available from inside the DS6000 Information Center.

Table 9. LED status after successful installation

LED Icon	LED Name	Appropriate state
	Power-on	Solid Green
	Location	Solid Blue

Table 9. LED status after successful installation (continued)

LED Icon	LED Name	Appropriate state
	Information	Unlit
	System alert	Unlit

- If all LEDs are not displaying the appropriate state, see the light path information to diagnose the problem.

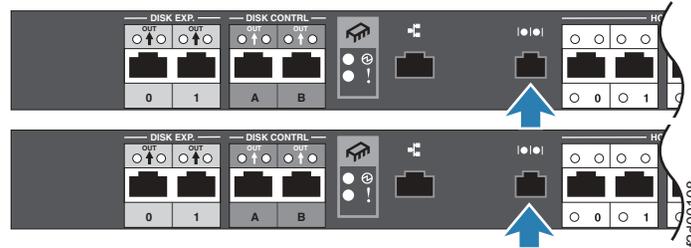
Setting the IP address

Complete this task to set the IP addresses and other network information for the processor cards.

You can create a direct connection from your computer to either of the processor cards on the DS6000. You can use the direct connection to set IP addresses, network masks, gateway IP addresses, and other network settings for one or both of the processor cards.

To set the IP address, perform the following steps:

- Use the serial conversion cable that was shipped with the enclosure to attach your computer to the serial port on the enclosure.



- Use a terminal emulator, such as NetTerm or Windows[®] HyperTerminal to connect to the server enclosure through the serial port that is located on the processor card. If you are using Windows HyperTerminal, you might need to provide a connection name and icon before you can specify the connection settings.

Choose a communications port, such as COM1 or COM3, and connect to the processor card using the following settings:

Remote connection setting	Remote connection value
Bits per second	38400
Data bits	8
Parity	None
Stop bits	1
Flow control	Hardware

- Use the default user ID **guest** and the default password **guest** to access the processor card. Press Enter to get the command prompt.
- If this is your first time setting the IP address, change the default **guest** password to one that you select.

- a. Select Change 'guest' password from the Main Menu options.
- b. Enter the current password. This is guest if you are changing it for the first time.
- c. Enter the new password that you have chosen. You will receive a message stating that the password has been changed successfully.

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.

Note: You must use the serial conversion cable to connect your computer to the second processor card to change the password on that processor card. You cannot change the password for the second processor card while you are still connected to the first processor card.

5. Select Configure network parameters from the Main Menu options.
6. Select Use static IP addresses from the Network configuration menu options. Set the IP addresses and network masks for both processor cards by performing the following steps:
 - a. Change the IP address for the processor card to which you are connected. Select IP address for this node from the Static IP addresses configuration menu options, enter the new IP address for the current processor card, and press Enter.
 - b. Change the IP address for the other processor card. Select IP address for other node from the Static IP addresses configuration menu options, enter the new IP address for the other processor card, and press Enter.
 - c. Change the network mask for the processor card to which you are connected. Choose Network mask for this node from the Static IP addresses configuration menu options, enter the new network mask for the current processor card, and press Enter.
 - d. Change the network mask for the other processor card. Choose Network mask for other node from the Static IP addresses configuration menu options, enter the new network mask for the other processor card, and press Enter.
 - e. Select Gateway for this node from the Static IP addresses configuration menu options, enter the new Gateway IP address for the current processor card, and press Enter.

Note: To assign a storage unit without a gateway, enter the address 0.0.0.0.
 - f. Select Gateway for other node from the Static IP addresses configuration menu options, enter the new Gateway IP address for the other processor card, and press Enter.

Note: To assign a storage unit without a gateway, enter the address 0.0.0.0.
 - g. Select Back to Network Configuration to return to the Network configuration menu.
7. Select Advanced Configuration Options to set the domain name server and the gateway settings.

- a. Select First DNS server from the Advanced network configuration options, enter the new DNS server address for the current processor card, and press Enter.
 - b. Select Second DNS server from the Advanced network configuration options, enter the new DNS server address for the other processor card, and press Enter.
 - c. Select Back to Network Configuration.
8. Select Back to Main Menu to return to the Main Menu.
 9. Select Apply changes and exit from the options in the main menu to save your changes and exit the application. You can then disconnect the serial conversion cable from your computer to the processor card.

Ensure that you set the IP address on the management console and that it is cabled to the same subnet as the DS6000. Change the default **guest** password for the second processor card. This processor card might appear to already be set, but you must still change the default **guest** password.

Enabling multiple IP addresses on the DS6000 Storage Manager

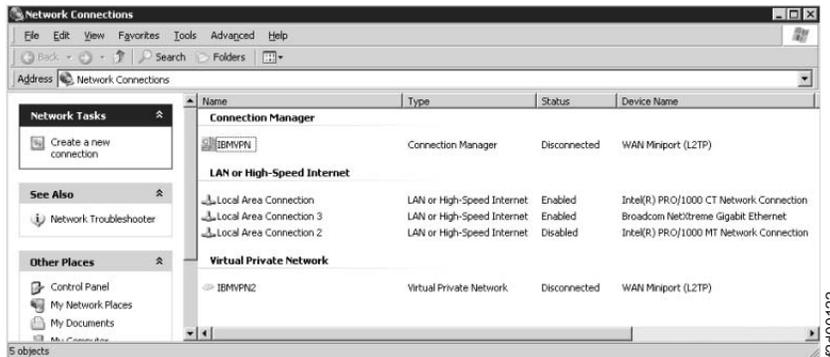
Complete this task to allow multiple IP addresses on the Storage Manager.

If you have multiple IP address on the DS6000 Storage Manager, ensure that the first network adapter is the adapter that is on the same subnet network as the DS6000. If this is not the case, the binding order must be changed so that the IP address on the same subnet (private) as the DS6000 is listed first in the binding order.

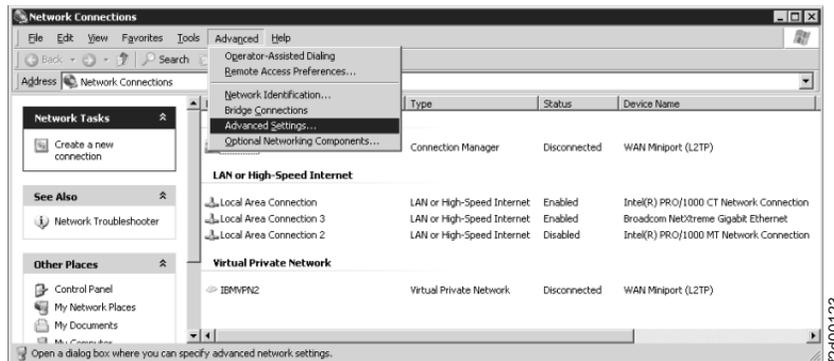
To change the binding order for the IP address, perform the following steps:

1. From the **Control Panel**, select **Network Connections** → **Open**.

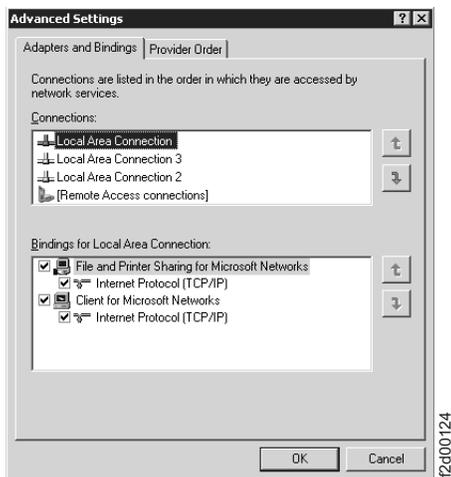




- From the **Network Connections** window, select **Advanced** → **Advanced Settings**.



- In the **Adapters and Bindings** tab, check the list of the network adapters to ensure that the first network adapter is the adapter that is on the same subnet network as the DS6000. If this is not the case, you must change the binding order so that the IP address on the same subnet (private) as the DS6000 is listed first in the binding order.



Installing a modem on the management console

Complete this task to install the modem on the management console.

You must have prepared a dedicated analog telephone line before you install the modem.

You must use this modem or a network VPN connection during remote support situations. A modem must be installed on the management console to enable the Call Home feature. You can use the modem to initiate a virtual private network (VPN) connection to IBM during remote support situations. However, if one or more phone numbers are stored on this page, remote support connections cannot initiate a VPN connection over the Internet. You must delete all phone numbers on this page before you can initiate an Internet VPN connection.

To install the modem on the management console, perform the following steps:

1. Remove the modem and accompanying cables from the package. Save the CD and any documentation that came with the modem for future use.
 2. Perform the following steps to connect the power cable from the modem to the power outlet in the wall:
 - a. Connect the power cable to the power adapter.
 - b. Connect the power adapter to the modem.
 - c. Connect the ac power cable to the wall power outlet.
 3. Connect the analog telephone line into the port on the modem that is labeled, "line".
 4. Plug the other end of the analog telephone line into the appropriate telephone outlet.
 5. Connect the modem to the management console. You can connect the modem to the management console in one of two ways:
 - Connect the modem to the management console directly through a serial port on the management console. The 9-pin serial connector connects to the management console while the 25-pin serial connector connects to the modem.
 - Connect the modem to the management console through the USB port on the management console. The USB-to-serial-port converter connects to the USB port on the management console while the 9-pin serial connector connects to the USB-to-serial-port converter and the 25-pin serial connector connects to the modem.
- Note:** Use the instructions that were provided with the USB-to-serial-port converter to install the converter.
6. Ensure that the management console is turned on.
 7. Use the front power switch to turn on the modem.
 8. Perform the following steps to install the modem on the management console:
 - a. Open the Phone and Modem Options settings for your operating system. In Windows 2000, Windows 2003, and Windows XP, click **Start** → **Settings** → **Control Panel** → **Phone and Modem Options**.
 - b. In the Dialing Rules tab, either click **New** to create a new dialing location, or click **Edit** to modify a previously selected dialing location.
 - c. Move through the General, Area Codes, and Calling Card tabs and enter the required dialing information for your location, and then click **OK**.
 - d. Select the Modems tab.
 - e. Select **Add**.
 - f. In the Install New Modem panel, ensure that the **Don't detect my modem; I will select it from a list** box is checked, and then click **Next**.
 - g. Click the **Have Disk...** button.

- h. Insert the CD that came with the modem into the CD-ROM drive on your management console. On the Install From Disk panel, navigate to the CD-ROM directory, and then click **OK**.
 - i. Select **MultiTech MT5600BA V9** from the Models list, and then click **Next**.
 - j. Select the communications port where you installed the modem, and then click **Next**. If a hardware error indicates that compatibility testing did not pass, click **Yes** to continue.
 - k. When a message appears that confirms that the installation is successful, click **Finish**.
 - l. Click **OK** to close the Phone and Modems Options window.
9. Configure modem phone numbers.
- a. In the navigation, under Real-time manager, select **Manage hardware** and then select **Storage complexes**. Select **Configure Modem Remote Support** in the **Select Actions** list and click **Go**.
 - b. Select the appropriate country. If applicable, select the appropriate state. A list of appropriate phone numbers is displayed.
 - c. Select the phone number that you want to use.
 - d. If necessary, specify the prefix that the modem must dial to reach an external phone line.
 - e. Click the **Populate** button next to an empty Phone# field. This places the phone number that you selected into the empty field.
 - f. If necessary, manually modify the phone number to work with dialing protocol for your location.
 - g. Click the **Test** button next to the phone number. This creates a modem connection.
 - h. Verify that the connection was created successfully. You can use the Test connection status field to view the current state of the connection.
 - i. Click the **Cancel test** button to end the test connection before the test is completed.
 - j. Select another phone number from the available list and use the **Populate** button to add additional phone numbers.
 - k. Click **OK** to store the modem phone numbers. Click **Cancel** to exit the page without saving any changes.

Chapter 4. Installing the DS6000 Storage Manager

The DS Storage Manager is installed using graphical or unattended (silent) mode for Windows operating systems. It can be accessed remotely from any location that has network access using a Web browser. The result of the installation is access to a configuration tool that can be used in offline or online settings. In addition, you can access licensed functions, such as FlashCopy, that have been activated on your storage unit.

Ensure that you complete and have available the following planning work sheets. They can be found in the *IBM System Storage DS6000 Introduction and Planning Guide*, Chapters 5 and 6.

- Host attachment work sheet
- Activation codes work sheet

You can also use the table in “Gathering installation information” on page 29 to gather information that you will need during installation.

Perform the following steps to install the DS Storage Manager:

1. Install the DS Storage Manager from the CD-ROM that came with your DS6000 series. See “Installing the DS Storage Manager on the Windows operating system using the graphical mode” on page 76 for information about installing the DS Storage Manager in graphical mode or “Installing the IBM System Storage DS Storage Manager on the Windows operating system in unattended (silent) mode” on page 82 for installing in unattended mode.
2. Restart your management console.
3. Start the IBM System Storage DS Storage Manager. See “Starting the DS Storage Manager on a Windows operating system” on page 89 for information on how to start the DS Storage Manager.
4. Set up the DS Storage Manager for use with your DS6000. See “Setting up the DS Storage Manager” on page 90 for information about setting up the DS Storage Manager.

Terminal Server:

Use these procedures if your Windows PC has the Terminal Server enabled and the session that you are logged in to is set to the “Execute” mode. You must change the mode to “Install” using the following instructions:

- For Windows 2000: Use the **Add New Programs** option of the Add/Remove Programs utility. This utility automatically changes the Terminal Server session to “Install” mode before you start the installation, and changes it back to the initial mode after you finish the installation.
- For Windows 2003: Double-click on the executable file that has the standard name of `setup.exe` to automatically change the Terminal Server session to “Install” mode before you start the installation. Windows considers `setup.exe` as an installation program. This utility automatically changes it back to the initial mode after you finish the installation.

Note: This behavior applies only to double-clicking on the `setup.exe`. It does not apply to invoking the `setup.exe` from a command line or launching it from an executable of a script file.

Internet browser support

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

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

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

Note: 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**.

Management console requirements

The DS6000 Storage Manager requires that the system that is used as the management console be continuously available for customer operation, configuration, and problem management.

It is best to load the DS6000 Storage Manager on a dedicated system. If you must install the DS6000 Storage Manager on a nondedicated system, contact your account sales representative to contract locally delivered services to install and set up the DS6000 Storage Manager. Do not use a laptop for the management console, because communication with the DS6000 is lost if the laptop is shut down or moved.

Before you can install the DS Storage Manager software on the computer that you will use to manage your storage system, you must ensure that your computer meets a minimum set of hardware and operating system compatibility requirements:

- IBM Personal Computer compatible with the following components:
 - 1.4 GHz Pentium® 4 processor
 - 256 KB cache
 - 1 GB memory
 - 1 GB disk space for the DS Storage Manager software
 - 1 GB work space per managed server enclosure
 - Ethernet IP connectivity to each processor card
 - Ethernet IP connectivity to external network (for call home and remote support)
 - USB port or serial port for modem connection
 - CD connectivity

If you want to order a management console, consider the IBM 8143 ThinkCentre™ M51 Model 34U (8143-34U) Desktop system with a 3.2 GHz/800 MHz Intel® Pentium 4 Processor. If a monitor also is required, IBM suggests the IBM 6737 ThinkVision C170 (6737-P6N) 17-inch full flat shadow mask CRT color monitor with a flat screen (16-inch viewable-image).

Ensure that your Windows PC is using one of the supported operating systems that are listed in Table 10.

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

Table 10. Supported operating systems

Operating System	Full management console install
Windows Server 2003 Enterprise Edition	X
Windows Server 2003 Standard Edition	X
Windows 2000 Advanced Server SP4	X (English only) See Note.
Windows 2000 Server SP4	X (English only) See Note.
Windows 2000 Professional SP4	X (English only) See Note.
Windows XP Professional SP1	
Windows XP Professional SP1a	
Windows XP Professional SP2	X
Note: Windows 2000 Advanced Server SP4, Windows 2000 Server SP4, and Windows 2000 Professional SP4, require update 818043.	

Ensure that your Windows PC has the regional settings option that is specified to your language and country or region.

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

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

To allow your browser to display the DS Storage Manager and the Information Center in separate windows, ensure that the Internet options are not set to reuse windows for launching shortcuts.

1. Select **Start, Settings, Control panel**, then **Internet Options**.
2. On the **Advanced** tab, ensure the **Reuse windows for launching shortcuts** checkbox is not selected.
3. Click **Apply**.

Installing the DS Storage Manager on the Windows operating system

This section contains information for DS Storage Manager installation on the Windows operating system.

The DS Storage Manager is installed using a graphical or unattended (silent) mode for the Windows operating systems. It can be accessed remotely using a Web browser from any location that has network access.

You can choose to install the IBM System Storage DS Storage Manager on the Windows operating system using either of the following modes:

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

After you have installed the DS 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 DS Storage Manager application, which includes the real-time and simulated manager components. These components are designed to help you create and manage the physical and logical configurations of your storage complexes and storage units. Plus, the real-time manager application provides you the opportunity to use the Copy Services features that you have purchased.

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

Complete this task to install the IBM System Storage DS Storage Manager in your Windows environment using the graphical mode.

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 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 DS6000 Series 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:

DS6000 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.

DS6000 Storage Manager Web site

Provides information from the product Web site.

Installation wizard

Starts the IBM System Storage DS Storage Manager installation program.

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. Click **Next**. One of the following is displayed:
 - The Server Parameters window is displayed with default values shown for all the fields.
 - If the DS Storage Manager Server is already installed on your system, the DS Storage Manager Server Installation Checking window is displayed. 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, click **Next**.
When you click **Next**, if the DS Network Interface Server is installed on the system, the DS Network Interface Server Installation Checking window appears. This 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, click **Next**.
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 15 on page 81.
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 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 recommend that the port numbers specified in this section not be changed from the defaults. Doing so, may 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 DS6000/Console; for example, `http://myhost.com:8451/DS6000/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 may 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. You can find details for the post installation tasks in "Setting up the DS Storage Manager" on page 90. 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\DS6000 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. If a restart is not required, exit the LaunchPad program by clicking **Exit** on the LaunchPad window.
 18. Verify that the DS Storage Manager server and IBM System Storage DS Network server are operating. From the Windows Control Panel, select Administrative Tools, then Services. Both servers should have a "Started" status. You cannot start the DS Storage Manager until after both of the servers are started.
 19. Start the DS Storage Manager. See "Starting the DS Storage Manager on a Windows operating system" on page 89.
 20. Complete the postinstallation tasks. You can find details for the postinstallation tasks in "Setting up the DS Storage Manager" on page 90.

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

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

Note: Skip this topic if you have already installed the DS Storage Manager on the Windows operating system using the graphical mode.

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.

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.

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: Tue Jan 25 18:01:00 EET 2005
# 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.
#####

-silent

#####
#
# IBM System Storage DS6000 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="C:\Program Files\IBM\DS6000StorageManager"

#####
#
# User Input Field - type
#
# The management types selection, can be: "full", which installs the Full
# Management Console, and "offline", which installs the Offline Management
# Console.
#

-W managementType.type="full"

#####
#
# User Input Field - hostname
#
# The fully-qualified host name of the machine where the 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.
# 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="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. To enable HTTPS, the user must also perform the
# procedure described in Console Developer InfoCenter (Setting up SSL) after DS
# Storage Manager Server is installed.
#

-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 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="8455"

#####
#
# 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="generate"

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

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

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

-W wasSslConfiguration.keystorePassword=""

#####
#
# User Input Field - confirmedKeystorePassword
#
# The password confirmation of the keystore file. On silent mode the password
# confirmation is NOT necessary.
#

-W wasSslConfiguration.confirmedKeystorePassword=""

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

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

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

-W wasSslConfiguration.truststorePassword=""

#####
#
# User Input Field - confirmedTruststorePassword
#
# The password confirmation of the truststore file. On silent mode the password
# confirmation is NOT necessary.
#

```

```

-W wasSslConfiguration.confirmedTruststorePassword=""

#####
#
# 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"

```

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.

Note: The DS Storage Manager 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.

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 DS6000 Storage Manager** and then click **Open DS Storage Manager**. The DS6000 Storage Manager is opened in the browser that you have set as your default.
4. If you are prompted to install the security SSL certificate. Click **Yes**. The IBM System Storage DS6000 Signon window is displayed when the DS6000 Storage Manager starts.

Note: If a firewall exists between the management console and the storage unit, ports 1700 through 1799 must be enabled. If the ports are not enabled, communication between the management console and the storage unit will fail.

5. 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.

Setting up the DS Storage Manager

Complete this task to complete the initial steps that are required after you install the DS Storage Manager.

You must have completed installing the DS Storage Manager and have restarted your management console before you can complete these steps.

If you are working in simulated mode, you must have completed each of the following subtasks before you can complete this task:

- Create a configuration file using simulated mode.
- Create the storage unit using simulated mode.
- Create a storage complex using simulated mode.

This task contains high-level steps that you must complete before you can create an initial logical storage configuration. Each of these steps contains links to individual sets of substeps. If you are completing this process in simulated mode, you must return and complete the real-time steps on the management console that is connected to your server enclosure before you can apply your simulated configuration.

Return to the next step of this task after completing each of the subtasks below.

1. Assign the storage unit to a storage complex. To assign the storage unit to a storage complex, you must complete the subtask in “Assigning a storage unit to a storage complex (real-time only)” on page 94.
2. Verify, and if necessary, upgrade the code level on the DS6000. You must complete the subtasks in “Verifying the current code level” on page 95 and Chapter 9, “Upgrading the DS6000 code level,” on page 199 to check your current level of code and apply any available updates.
3. Activate your licensed functions. You must complete the subtasks in “Activating licensed functions” on page 95. This includes obtaining your feature activation codes, importing your feature activation codes, and applying the feature activation codes to your storage unit.
4. (Optional) Define a peer management console. To define a peer management console, you must complete the subtask in “Defining multiple management consoles (real-time only)” on page 97.
5. Specify the storage unit day and time. To specify the storage unit day and time, you must complete the subtask in “Specifying storage unit day and time (real-time only)” on page 98.
6. Specify the storage unit network settings. To specify the storage unit network settings, you must complete the subtask in “Reviewing storage unit network settings (real-time only)” on page 99.
7. Define customer contact information. To define customer contact information, you must complete the subtask in “Defining customer contacts (real-time only)” on page 99.
8. Configure your Call Home, SNMP, and zSeries Service Information Message (SIM) notifications. To configure your notifications, you must complete the subtask in “Setting up Call Home, SNMP, and SIM notifications (real-time only)” on page 100.
9. Configure the I/O ports. To configure the I/O ports, you must complete the subtask in “Configuring I/O ports” on page 101.
10. Create the host systems. To create the host system, you must complete the subtask in “Creating host systems” on page 102.

Creating a configuration file (simulated only)

Complete this task to create a simulated configuration file.

A simulated configuration file allows you to store management console and logical storage configuration settings to be applied to the DS6000 at a later time. The

configuration file contains information for one or more storage units (including both physical and logical) and one or more host systems.

1. In the navigation, select **Simulated manager** → **Manage configuration files**. From the **Select Action** menu, select **Create new...**, and then click **Go** to create a new enterprise file for offline configuration. A default simulated configuration file is available to use if you do not want to create a new file.
2. If you have another file open when you select the **Create new...** action, a message prompts you to save your current work before you create a new file. If you click **Ok** after you receive the message, the new configuration file opens after it is created. If you click **Continue**, you must open the newly created configuration file before you can start working in it.

Once you create a simulated configuration file, you must either import a storage unit instance or create a new storage unit instance from the Create storage unit page.

The configuration file that you create has a default name until you save it using the **Save** action. The default name for the enterprise files is “Enterprise 1,” incremented by one for each existing default file name that you do not change.

If you are creating a new logical storage configuration, return to “Creating a custom logical storage configuration” on page 167 and complete the rest of the steps.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Creating a storage unit (simulated only)

Complete this task to create a simulated storage unit and to specify its attributes and properties.

1. In the navigation, select **Simulated manager** → **Manage hardware** → **Storage units**.
2. From the **Select Action** menu, select **Create...**, and then click **Go**. The Create Storage Unit — General storage unit information page is displayed.
3. In the General storage unit information page, you must specify the machine type and nickname. The other fields are optional. You can enter the **Select storage complex** value now or modify the storage unit properties later. If you want to create a new storage complex, click the **Create new storage complex** button. The new complex is listed for your selection after you finish the creation wizard process.
4. Click **Next** to continue.
5. The Create Storage Unit — Specify DDM packs page is displayed. You must specify the **Quantity of DDM packs** and the **DDM type**. Click **Add**, and then click **Next** to continue. The Define licensed function page is displayed.
6. In the Define licensed function page, you must specify a value in the **Operating Environment License (TB)** field. You can optionally specify values in the remaining four fields as appropriate. After you enter the necessary values, click **Next** to continue.

The **Operating Environment License (TB)** value is the total amount of capacity in the box. If you specify more than one storage unit, the license is split equally between the two storage units.

7. If you specified a license for the FlashCopy or Remote Mirror and Copy functions, you must specify the storage type (FB, CKD, or All types) for each of the Copy Services functions. Once you have defined the storage type, click **Next**.
8. The Verification page is displayed. Use this page to review the established attributes and verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and specify the correct values. Otherwise, click **Finish** to complete the storage unit creation process.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Creating a storage complex (simulated only)

Complete this task to create a simulated storage complex and to specify its nickname and storage unit assignments.

1. In the navigation, select **Simulated manager** → **Manage hardware** → **Storage complexes**.
2. From the **Select Action** menu, select **Create...**, and then click **Go**. The Create Storage Complex — Define properties page is displayed.
3. If a storage unit is not already defined, you must specify a nickname, which is limited to 16 characters. The other fields are optional.
 - a. The **Available and Selected Storage units** fields are not required when you create the storage complex. You can enter this value now or modify the storage complex properties later. Additionally, you can select this storage complex when you create a storage unit. You must, however, create an association between a storage complex and a storage unit at some point before you download or upload configurations to or from the storage unit.
 - b. If you select the **Create new storage unit** button, the new storage unit is available for selection after you complete the create storage unit process.
4. After you have defined the properties, click **Next** to continue. The Verification page is displayed.
5. Use the Verification page to review the established attributes and verify that they are correct. If the attributes and values are not correct, click **Back** to return and specify the correct values. Otherwise, click **Finish** to complete the storage complex creation process.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Locating the serial number, model, and signature

You will need to record your machine’s serial number, model, and signature so that this information is available when you are required to enter it to obtain your feature activation codes.

To locate the serial number, model, and signature, perform the following steps:

1. Locate the MTMS label on the flange on the front right side of the enclosure at the right front bezel. The MTMS label is a black label with white lettering.
2. From the MTMS label, record the model and serial number in the table below. The Machine Type - Model Number - Serial Number (MTMS) is a string that contains the machine type, model number, and serial number. Only the last

seven characters of the string are the machine serial number. For example, if the MTMS is IBM.1750.511.75FA120, then the machine type is 1750, the model is 511, and the serial number is 75FA120.

3. Start the DS Storage Manager application.
4. In the navigation panel select, **Real-time Manager** → **Manage Hardware** → **Storage Units**. On the storage unit main page, select the storage unit, click **Properties** in the **Select Action** list, and then click **Go**. The properties page displays for the storage unit.
5. From the **Machine signature** column, note the machine signature in the table below.

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

Assigning a storage unit to a storage complex (real-time only)

Complete this task to assign a storage unit to the selected storage complex and specify the appropriate network settings.

This process must be done from the primary management console. You must make a selection in the table to enable this option.

To complete this task, you must know your machine's serial number. See "Locating the serial number, model, and signature" on page 93.

Note: Before you configure, disable any firewalls, as they might interfere with DS6000 communication.

1. In the navigation, select **Real-time manager** → **Manage hardware** → **Storage complexes**. On the storage complexes main page, select the appropriate storage complex from the table.
2. From the **Select Action** menu, select **Assign Storage Unit**, and then click **Go**. The Assign Storage unit — Storage unit properties page is displayed.
 - a. Enter a Nickname.
 - b. Optionally, enter a description.
 - c. Enter the IP address of processor cards 1 and 2.
 - d. The machine type is already generated, but you must enter the 7-digit serial number, without hyphens, from the MTMS label on the flange on the front right side of the enclosure at the right front bezel. See "Locating the serial number, model, and signature" on page 93.
3. Click **Next**. The Network settings page is displayed.
4. Specify the appropriate network settings and then click **Next**. The Verification page is displayed.
 - a. Enter a gateway.
 - b. Enter a subnet mask.
 - c. Enter the primary DNS address.
 - d. Enter the secondary DNS address.
 - e. If necessary, enter a different Max transmission units value. This is the maximum rate for transmission. The valid range is 1 to 9000 bytes. The default is 1500 bytes, which is appropriate in most cases.

5. Verify the attributes and values for the newly configured Storage unit. Click **Finish** if the settings are correct.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Verifying the current code level

Complete this task to verify the current level of code on your DS6000.

1. In the navigation, select **Real-time manager** → **Manage hardware** → **Storage units**. On the Storage units - Main Page, select the appropriate storage unit.
2. From the **Select Action** list, select **Apply firmware update**, and then click **Go**. The Apply firmware update page is displayed.
3. Locate the Current firmware level row in the table to view your current level of code.
4. Compare your current code level with the latest available code level. Follow the instructions in Chapter 9, “Upgrading the DS6000 code level,” on page 199.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Activating licensed functions

This section contains information to help you 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 DS6000 licensed function is a concurrent activity (assuming the appropriate level of microcode is installed on the machine for the given function).

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 order confirmation code when you enter the DS6000 (1750) serial number and signature. However, if the order confirmation code is not attached to the 1750 record, you must assign it to the 1750 record in the DSFA application. In this situation, you will need the order confirmation code (which you can find on the License Function Authorization document).

Importing activation codes

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

Notes:

1. Before you begin this task, you must resolve any current open problems. See the Troubleshooting section of the IBM System Storage DS6000 Information Center for additional information. If you need additional assistance to resolve these problems, contact IBM Support.
2. Before you configure, disable or provide paths through any firewalls, because they might interfere with DS6000 communication.

Perform the following steps to import your activation codes:

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 import the activation code for.
2. From the **Select Action** menu, select **Configure**, and then click **Go**. Select the **Activation codes** tab. The Apply activation codes 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.

Applying activation codes

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

Notes:

1. The initial enablement of any optional DS6000 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 initialization or reboot:
 - Removal of a DS6000 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 resolve any current open problems.
3. Before you configure, disable or provide paths through any firewalls, because they might interfere with DS6000 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 units**. On the Storage units main page, select the storage unit that you want to import the activation code for.
3. From the **Select Action** menu, select **Configure...**, and then click **Go**. The Storage unit properties page is displayed. Click on the **Activation codes** tab. The Activation codes page is displayed.
 - a. If you already imported your activation codes from a file or retrieved existing codes from the storage unit, the values are displayed in the fields and you can modify or overwrite them, as appropriate.
 - b. If you are importing your activation codes from a file that you downloaded from the DSFA Web site, click **Import key file**. Once you complete the import process, the data from the file is displayed.
 - c. If you did not download your activation codes into a file, enter the codes into the appropriate fields.

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

4. Click **OK** to complete the process.

Defining multiple management consoles (real-time only)

Complete this task to create a storage complex domain by establishing a connection with a secondary management console (peer) for redundancy.

You can install a secondary management console (MC) for backup purposes. It is especially important if you use DS Copy Services. You install the secondary MC the same way that you installed the primary MC.

After the primary and secondary MCs are installed, perform this process from the primary MC. You must have the MC IP address and the appropriate user ID and password. To establish a peer-to-peer relationship, both user IDs and passwords must be the same for the primary and secondary MCs.

When you add a secondary MC, the secondary MC cannot have storage units that are associated with it. If you have two MCs that each have storage units that are associated with them, you must remove the storage units from the secondary MC that you want to use as the peer of the primary MC. After you remove the storage

units, add the secondary MC to the primary MC. You can then add the previously removed storage units using either the primary or the secondary MC.

You must have the same administrator password set for both MCs before you can define multiple management consoles.

Perform the following steps to establish a connection with a secondary MC for the storage complex that is associated with the primary MC.

Note: When you use peer MCs, if one of the consoles fails, you cannot make changes to any user accounts. If the peer-to-peer relationship cannot be restored, you must remove the secondary MC using the Removing Peer Management Consoles function, using the primary MC.

1. In the navigation, select **Real-time manager** → **Manage hardware** → **Storage complexes**. On the Storage complexes – Main page, select a storage complex.
2. From the **Select Action** menu, select **Define peer**, and then click **Go**. The Define peer management console page is displayed.
3. Perform one of the following steps, but not both:
 - Specify the server host name for the secondary MC.
 - Specify the server IP address for the secondary MC.
4. Click **OK**. The storage complex domain is established.
5. Open a command prompt and navigate to the C:\Program Files\IBM\dsniserver\bin\ directory where the recovery tool (script) has been installed.
6. Type the script name, `securityRecoveryUtility.bat -s`

Note: This utility synchronizes all of the user account information between the primary and secondary MCs. After the user accounts have been synchronized, all changes to user accounts are automatically reflected on both peer MCs unless the peer-to-peer relationship is broken. If you remove a peer, and define a new peer, you must run the recovery tool script again after you determine the new peer.

Click on the **Storage Complexes** main page and you should see a second MC nickname that is defined with status on the right-most column.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Specifying storage unit day and time (real-time only)

Complete this task to specify date, time, time zone, and Daylight Saving time observation setting for the selected storage unit.

Important: The date and time on the storage unit must be set correctly and must match the date and time settings on the management console. If this information is not set correctly, the error logs for the storage unit will not be correct.

You must make a selection in the table to enable this option.

To specify storage unit day and time, perform the following steps:

1. In the navigation, select **Real-time manager** → **Manage hardware** → **Storage units**. Select the appropriate storage unit.

2. In the **Select Action** list, select **Configure**, and then click **Go**. The Storage unit properties page is displayed.
3. In the navigation on the left, click the **Date and time** tab. The Date and time zone page is displayed.
4. Specify the date, time, and time zone for the selected storage unit.
5. Click **OK** to save and close.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Reviewing storage unit network settings (real-time only)

Complete this task to view properties for the selected storage unit and optionally modify the nickname and description.

You must make a selection in the table to enable this option.

To review storage unit network settings, perform the following tasks:

1. In the navigation, select **Real-time manager** → **Manage hardware** → **Storage units**. On the Storage units — Main page, select the appropriate storage unit.
2. From the **Select Action** menu, select **Configure**, and then click **Go**. The Storage unit properties page is displayed.
3. In the navigation on the left, click the **Network settings** tab. The Network settings page is displayed.
4. Review the IP addresses and host names for the selected storage unit.
5. Optionally modify the nickname and description.
6. Click **OK** to save and close.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Defining customer contacts (real-time only)

Complete this task to add or modify shipping or contact information for a customer account.

Important: When you configure the Call Home feature, you must provide complete, accurate, up-to-date contact information, including the correct telephone number. Incomplete or incorrect contact information can cause a delay in the IBM response to a Call Home event. If your contact information changes after you initially configure Call Home (for example, if you move the DS6000 to a new location) be sure to update your customer contact information.

To define customer contacts, perform the following steps:

1. In the navigation, select **Real-time manager** or **Simulated manager** → **Manage hardware** → **Storage units**. On the Storage units — Main page, select the appropriate storage unit.
2. From the **Select Action** list, select **Customer contact**, and then click **Go**. The Customer account information tab is displayed.
3. Create or modify the customer account information. Required fields are indicated by an asterisk (*). Fill in all required fields or you will be unable to set up the SMTP server. You are required to enter information in the Business/company name field.

4. Click the Contact information tab and add or modify contact information for the customer.
5. Click the Shipping information tab and add or modify shipping information for the customer.
6. Click the **OK** button to complete the customer contact information.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Setting up Call Home, SNMP, and SIM notifications (real-time only)

Complete this task to configure your Call Home (SMTP or modem), SNMP, and zSeries Service Information Message (SIM) notifications.

You must define your customer contact information before configuring notifications, or you will be unable to set up the SMTP server.

The Call Home feature allows the transmission of operational and error-related data to IBM. It provides the ability for the storage unit to alert IBM support to machine conditions. The SNMP feature generates alert messages and sends them to your designated location.

Important: When you configure the Call Home feature, you must provide complete, accurate, up-to-date contact information, including the correct telephone number. Incomplete or incorrect contact information can cause a delay in the IBM response to a Call Home event. If your contact information changes after you initially configure Call Home (for example, if you move the DS6000 to a new location) be sure to update your customer contact information.

Notes:

1. The management console does not process SNMP messages. You are responsible for installing additional SNMP alert software to process these messages.
2. You must enable your SMTP server to relay SMTP messages from your DS6000.

To define Call Home (SMTP or modem), SNMP, and zSeries SIM (service information message) notifications for a storage unit, perform the following steps:

1. In the navigation, select **Real-time manager** → **Manage hardware** → **Storage units**.
2. From the **Select Action** menu, select **Configure notifications**, and then click **Go**. The Configure notifications — Define Call Home page is displayed.
3. Ensure that **Enable Call Home** is selected to activate the Call Home function. (This is checked by default.) You can also optionally check the **Enable Modem Call Home** box if you want to send Call Home notifications through your modem connection. You must have a modem installed and configured if you want to use this function. If both the **Enable Call Home** and **Enable Modem Call Home** boxes are selected, Call Home notifications are first sent through the modem connection and then through the SMTP connection if the modem connection is unsuccessful.
4. Complete the following SMTP information:
 - a. Enter the SMTP server host name (for example, server.company.com).

- b. Enter the SMTP IP address (for example, 123.456.7.89). This must be an IP address that the storage unit can reach.
 - c. Enter the SMTP server port (for example, 25). It is recommended that this port not be changed from the default port (25). Changing the port number can cause the Call Home feature to not work.
 - d. Click **Apply**.
5. Click **Test Call Home connection** to send a connection test and generate a problem log entry. A confirmation message is displayed. Click **Ok**.
 6. Click the **SNMP** tab. The Define SNMP connection page is displayed.
 7. Select Enable SNMP notification to define the SNMP connection properties for the selected storage units.
 8. Specify either an IP address, a Host name, or both under SNMP trap destination.
 9. Specify an SNMP community name of up to 32 characters. This field is used to authenticate requests. 'public' is selected by default.
 10. (Optional) Specify an SNMP system contact name of up to 32 characters.
 11. Enter a destination port.
 12. Click **Apply**.
 13. Click the **zSeries** tab. The Define SIMs for zSeries page is displayed.
 14. (Optional) Select a SIM severity level in the Severity reporting level for DASD Service Information Messages (SIMs) field.
 - **Acute:** An irrecoverable error with possible loss of data. Use this severity level only for DASD SIMs.
 - **Serious:** An irrecoverable error or a data check with loss of access to data.
 - **Moderate:** A system path is not operational and performance might be degraded. This severity level does not apply to media SIMs.
 - **Service:** A recoverable error, equipment checks, or data checks. You can defer repair.
 - **None:** No messages will be sent.
 15. (Optional) Select the Maximum number of additional times a DASD Service Information Message is to be sent [0-5].
 16. (Optional) Select a Media Service Information severity level in the Severity reporting level for Media Service Information Messages field.
 17. (Optional) Select the Maximum number of additional times a Media Service Information Message is to be sent [0-5].
 18. (Optional) Select a Service Information severity level in the Severity reporting level for a storage unit Service Information Messages field.
 19. (Optional) Select the Maximum number of additional times a storage unit Service Information Message is to be sent [0-5].
 20. Click **Ok**.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Configuring I/O ports

Complete this task to change the configuration for I/O ports that have host attachments assigned to them.

1. In the navigation, select **Real-time manager or Simulated manager** → **Manage hardware** → **Storage units**. On the Storage units — Main Page, select a Storage unit.
2. From the **Select Action** menu, select **Configure I/O Ports...**, and then click **Go**. The Configure I/O Ports page is displayed.
3. Use the check boxes to select one or more host attachments of the same type.
4. From the **Select Action** menu, select the I/O port type that you want to change to. You can change any I/O port to FcAl , FcSf , or FICON.

FICON

Select this if you plan to connect one or more S/390 hosts using a FICON S/390 channel to the DS6000 either directly or through a FICON switch.

FcSf Select this if you plan to connect one or more open systems hosts to the DS6000 through a fabric switch, or if you plan to connect one or more S/390 hosts running LINUX on an FCP S/390 channel.

FcAl Select this if you plan to connect an open systems host directly to the DS6000 without going through a fabric switch.

5. Click **Go**. The table will update with the attachment type that you selected.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Creating host systems

Complete this task to create host systems and define their parameters.

1. In the navigation, select **Real-time manager or Simulated manager** → **Manage hardware** → **Host systems**. In Host Systems — Main Page, select a storage complex (and possibly a storage unit). From the **Select Action** menu, select **Create...** Then click **Go**. The Create Host System — General host information page is displayed.
2. In the General host information page, specify the host type and nickname and optionally provide a description. Then click **Next**. If you specified an open systems host, the Create Host System — Define host ports page is displayed; go to the next step.
3. In the Define host ports page, you must specify the quantity and attachment port type and you must click **Add** to add at least one host port definition to the Defined host ports table. You can optionally check the **Group ports to share a common set of volumes** box, so the quantity of ports identified in the Quantity field becomes grouped together and treated as a single host attachment.
4. Select at least one host port from the Define host ports table, and then click **Next**. The Create Host System — Define Host WWPN page is displayed.
5. In the Define Host WWPN page, specify the host port WWPNs for open systems hosts. Then click **Next**. The Create Host Systems — Specify storage units page is displayed.
6. In the Select storage units page, select each storage unit you want to access from the server by selecting a it from the Available storage units list and clicking **Add**. Then click **Next**. If you select the **Create a Storage Unit** button (Simulated only), follow the process for creating the new storage unit. Once you have completed that process by clicking on the **Finish** button, the new storage unit is available for selection.
7. In the Create Host Systems — Specify storage unit parameters page, specify the parameter values. Select a host attachment ID. Select a volume group to which

the host attachment needs access. (You can optionally choose **Select volume group later** if you do not want to select the volume group now.) Choose a login option to determine the FC adapter ports that the host can access. You can loop through this page for each host attachment identifier by selecting the **Apply assignment** button to commit the current transaction. You can start from the top to select another identifier. If you select an existing host attachment identifier from the table, you can click the **Create a new group** button to create a new volume group for selection. If you decide that this host attachment can log in to **the following specific storage unit I/O ports**, you must specify the specific ports in the available storage unit I/O ports table. When you are finished in the Specify storage unit parameters page, click **Apply assignment**, then click **OK**. The Create Host Systems — Verification page is displayed.

Note: You must click **Apply assignment** with at least one host attachment to the storage image before you can proceed to the Create Host Systems — Verification page.

8. In the Verification page, review the attributes and values to verify that they are correct.
9. If the attributes and values are not correct, click **Back** as appropriate to return and to specify the correct values. Otherwise, click **Finish** to complete the host system creation process.

| If you are setting up the DS Storage Manager, return to “Setting up the DS Storage
| Manager” on page 90 and complete the rest of the steps.

Chapter 5. Installing the DS6000 Command-Line Interface

The DS Command-Line Interface (CLI) provides a full function command set that allows you to check your storage unit configuration and perform specific application functions when necessary. It also enables open systems hosts to invoke and manage FlashCopy, Metro and Global Mirror functions through batch processes and scripts.

You can install the DS CLI using a graphical, silent, or console mode. The mode that you use is determined by your needs and capabilities of your operating system.

Notes:

1. You cannot install the DS CLI on a Windows 64-bit operating system.
2. The installation process stops if the correct version of Java is not found. However, the DS CLI installation CD contains the Java 1.4.2 version. You can install this version on your system and continue with the installation of the DS CLI.
3. Ensure that you check the DS CLI README file on the installation CD. It contains the latest installation requirements that might not have made it into the installation documentation.
4. If you are using the console mode to install the DS CLI, you must open another console session to log into a DS CLI session when the installation is complete.
5. When the DS CLI installation process is complete, change to the directory where the application was installed and log in.

Preparing to install the DS CLI

This section contains information about preparing to install the IBM System Storage DS CLI.

You can install the DS CLI on a number of different operating systems. However, there are unique preparation instructions for the OpenVMS operating system.

Preparing for the DS CLI installation

The IBM System Storage DS CLI program enables open systems hosts to invoke and manage FlashCopy and Metro and Global Mirror functions through batch processes and scripts. You can choose to install the DS CLI interface using the unattended (silent) mode, console mode, or graphic mode. The mode of installation that you choose is often dependent on the configuration of your system.

Before you use the DS CLI commands, be aware of the following requirements:

- Your management console must be equipped with the DS Storage Manager graphical user interface (GUI).
- The GUI must have been installed as a Full Management Console installation.
- Your storage unit must be configured. The DS Storage Manager is used for this initial configuration. The configuration process includes the following tasks:
 - Selecting your storage complex
 - Assigning your storage unit to the storage complex

- Designating network information for the storage unit
- Before you can use the CLI commands that are associated with Copy Services functions, you must activate your license activation codes (part of DS Storage Manager postinstallation instructions).

The DS CLI can be installed on only one of the following supported operating systems:

Supported Host Systems	Installation File Name
IBM AIX (5.1, 5.2, 5.3)	setupaix.bin
Hewlett-Packard-UX (11.0, 11i, v1, v2)	setuphp1020.bin or setuphp11x.bin
Linux (Red Hat 3.0 Advanced Server [AS] and Enterprise Server [ES])	setuplinux.bin for xSeries setupgenericunix.sh for pSeries
SUSE Linux SLES 8, SLES 9, SUSE 8, SUSE 9	setuplinux.bin for xSeries setupgenericunix.sh for pSeries
Sun Solaris (7, 8, 9)	setsolarisSparc.bin
HP Tru64 (5.1, 5.1A)	setupgenericunix.sh
Novell Netware 6.5	setupwin32.exe
System i system i5/OS 5.3	Add -os400 at a command line after the name of the installation file. For example, setupaix.bin -os400
Windows 2000, Windows Datacenter, Windows 2003, and Windows XP	setupwin32.exe
UNIX users that do not have an X display	setupos.bin -console
<p>Note: Use the same installation file name for your host operating system as shown in the previous rows, but add the -console parameter after the installation file name. For example:</p> <pre>setupaix.bin -console</pre> <pre>setuplinux.bin -console</pre>	Where <i>os</i> represents the name of your operating system as shown in the previous rows.

- The DS CLI cannot be installed on a Windows 64-bit system.
- You must have Java 1.4.1 or later installed on your machine. The installation program checks for this requirement during installation and does not install the DS CLI if you do not have Java 1.4.1 or later.

Note: The DS CLI installation CD-ROM contains Java 1.4.2, which you can install if your system is not equipped with this version of Java.

- You must use the ksh (Korn shell) or bash (Bourne again shell) shell if you are installing on a HP Tru64. Installshield does not support the sh shell. You must perform all HP Tru64 installations using the **setupgenericunix.sh** file that is located on the installation compact disc.
- AIX installation:
 - The LIBPATH environment variable can interfere with the installation of the DS CLI and can result in the display of the JVM-not-found error. To avoid this interference, disable the LIBPATH environment variable before you install the DS CLI. After the installation of the DS CLI, enable the LIBPATH environment variable so that it can be used with other applications.

- Run the following commands to sequentially disable the LIBPATH environment variable, install the DS CLI, and restore the LIBPATH environment variable:

```
export LIBSAVE=$LIBPATH
unset LIBPATH
setupaix.bin -is:javahome /opt/ibm-Java-whatever/java/bin
export LIBPATH=$LIBSAVE
unset LIBSAVE
```

- System i system installation:

Note: The installation of DS CLI on a System i system is done remotely from a Windows, AIX or some other system platform. You cannot install the DS CLI directly on a System i system.

The System i system and i5/OS must meet the following requirements before the DS CLI can be installed:

- Prerequisites
 - The latest Java group program temporary fixes (PTF)
 - i5/OS 5722-SS1 option 34 - Digital certificate manager
 - Licensed product 5722-AC3 option *base - Crypto Access Provider 128 bit
 - Licensed product 5722-DG1option *base - IBM HTTP Server for iSeries
 - Licensed product 5722-JV1 options 6 - Java Developer Kit 1.4
 - The latest cumulative PTF package that is installed on the i5/OS
- If you are installing onto a System i system, ensure that the machine that you are installing from is network-attached to the iSeries machine.
- During the installation of the DS CLI application onto a System i system, you must provide the following information:
 - The name of the iSeries machine to which you are installing the DS CLI application.
 - The user name and password that is used to access the designated iSeries machine.

The installation process installs the DS CLI in the following default directory:

AIX /opt/ibm/dscli

HPUX /opt/ibm/dscli

Linux /opt/ibm/dscli

Sun Solaris
/opt/ibm/dscli

Windows
C:\Program Files\IBM\dscli

HP Tru64
/opt/ibm/dscli

iSeries
/ibm/dscli

Novell Netware
SYS:\dscli

Perform the following steps in preparation for the DS CLI installation:

1. Log on to your host system as a root user or administrator.

2. Insert the DS CLI product CD into the CD drive. If a window opens for the CD drive, close the window.
3. Mount the CD drive using the **mount** command according to your system. You can mount your CD drive using the following examples:

AIX Create a directory for the CD-ROM by issuing the following command:
`mkdir /cdrom -p`

Create a file system for the CD-ROM by issuing the following command:

```
crfs -v cdrfs -p ro -d cd0 -m /cdrom
```

where *cd0* represents the CD-ROM drive.

Mount the CD-ROM file system by issuing the following command:

```
mount /cdrom
```

HPUX Mount the CD-ROM file system using the path name for your environment by issuing the following commands:

```
ioscan -funC disk | more  
mount /dev/dsk/c?t?d? /<cdrom>
```

Linux Issue the following command on Red Hat systems:

```
mount /dev/cdrom
```

Sun Solaris

Issue the following command:

```
mkdir /mnt  
mount -F hsfs -r /dev/dsk/c0t6d0s2 /mnt
```

Note: The device name `/dev/dsk/c0t6d0s2` is the default name for Sun Solaris. The device name might be different on your system depending on your hardware configuration.

Windows

You are not required to mount the CD if you are using this operating system.

HP Tru64

Issue the following command:

```
mount -t cdfs -o noversion /dev/rznn /mnt
```

where *nn* represents the number of CD-ROM drives.

Novell Netware

You are not required to mount the CD if you are using this operating system.

4. Navigate to your CD drive and proceed with either the unattended (silent), console, or graphic installation.

Preparing to install the DS CLI on an OpenVMS System

The IBM System Storage DS CLI program enables open systems hosts to invoke and manage FlashCopy and Metro and Global Mirror functions through batch processes and scripts. The DS CLI is installed on an OpenVMS system by using the Polycenter Software Installation utility.

Be aware of the following requirements before you begin the installation of the DS CLI:

- The installation process places all command-line interface files in the [IBMDSCLI...] directory tree as a subdirectory of the Polycenter destination directory. You can specify this directory by using the **PRODUCT INSTALL** command with the /DESTINATION=devicename:[directoryname] qualifier. If you omit the device name, the Polycenter installation utility uses your current default device. If you omit the directory name, the Polycenter Software Installation utility uses the [VMS\$COMMON] directory as the default destination directory.

If you do not use the /DESTINATION qualifier at all, the utility installs the software in the location that is defined by logical name **PCSI\$DESTINATION**. If this logical name is not defined, the utility installs the software in SYS\$SYSDEVICE:[VMS\$COMMON].

- Extract and check the Release Notes file from the DS CLI installation package in the root directory of the installation CD.
- Ensure that you have at least 140000 blocks (approximately 70 MB) of free space on the installation disk.
- Install the required prerequisite patches operating system patches (ECOs) on your OpenVMS host system.

The command-line interface installation process automatically installs Java™ Runtime Environment (JRE) 1.4.2-4 on your host system. The JRE requires several ECOs. For detailed patch information, see the Java SDK v1.4.2 patch installation page at:

http://h18012.www1.hp.com/java/download/ovms/1.4.2/sdk1.4.2_patches.html

To download these ECOs, use the HP IT Resource Center (ITRC) database to perform a search for the patches from which you can select. These download steps are described at the Java SDK v1.4.2 patch installation page.

Installing the DS CLI

This section contains instructions for installing the IBM System Storage DS CLI.

You can install the DS CLI on a number of different operating systems. However, there are unique installation requirements for the OpenVMS operating system.

You can install the DS CLI using one of the following methods:

- Silent mode
- Console mode
- Graphical mode

Installing the DS CLI using the unattended (silent) mode

Complete this task to install the DS CLI application using the unattended (silent) mode.

Consider the following before you perform the installation of the DS CLI application:

- You cannot install the DS CLI on a computer that contains the ESS CLI.
- You cannot install the DS CLI on a Windows 64-bit operating system.
- You must have installed a version of Java 1.4.1 or higher on your system in a standard directory. The DS CLI installer checks the standard directories to determine if a version of Java 1.4.1 or higher exists on your system. If this

version is not found in the standard directories, the installation fails. If the installation fails because the correct version of Java is not found, see “Correcting the JVM-not-found error” on page 125.

The DS CLI installation CD-ROM contains the Java 1.4.2 version. You can install this version from the CD and continue with the installation of the DS CLI application.

- If the DS CLI application was installed on your client/host machine in the past, ensure that you end any active sessions of the DS CLI before you run the DS CLI installation CD.
- If you are installing onto a Novell system, you are directed to provide the following information:
 - The directory where your Windows drive is mapped.
 - The directory where the JAVA HOME is located.
- If you are installing onto a System i system, ensure that the machine that you are installing from is network-attached to the iSeries machine.
- During the installation of the DS CLI application onto an System i system, you must provide the following information:
 - The name of the iSeries machine to which you are installing the DS CLI application.
 - The user name and password that is used to access the designated iSeries machine.

The unattended (silent) mode of installation allows you to perform the installation from the command line without prompts or feedback. You can create an options/response file and use a text editor to change the default selections to a selection of your choosing.

Before you can use the DS CLI, some operating systems require that you restart your system after the installation is complete. Or, you might be required to open a new command prompt window to start a DS CLI session.

Note: After you install the new version of DS CLI, your old DS CLI sessions might be unusable.

Perform the following steps to install the DS CLI application using the unattended (silent) mode:

1. Log on to your system as an administrator.
2. Insert the DS CLI installation CD into the CD drive. If the InstallShield starts, click the **Exit** selection to exit.
3. Type the following command for all systems except the System i system at the command prompt `setup<platform>.<exe | bin | sh> -silent`. Press the **Enter** key on your keyboard to start the installation process in unattended (silent) mode.

The silent installation process applies all the default options to your installation. If you want to modify the default options, go to the next step.

Note: Initially the DS CLI installer checks your standard directories for the correct version of Java. If the correct version of Java is not found, you receive the following message:

A suitable JVM could not be found. Please run the program again using the option `-is:javahome <JAVA HOME DIR>`

If you receive this message, see “Correcting the JVM-not-found error” on page 125.

- |
- |
- |
- |
- |
- |
- |
- |
- |
- |
4. Generate the template response file if you must modify the defaults. Two example response files are provided in the cliReadmes directory of the CD: one for Windows, Novell, and iSeries installations, and one for all UNIX and Linux installations. You can copy the applicable response file to a directory of your choice.

The following output is an example of the Windows, Novell, or iSeries response or options file.

```

InstallShield Options File Template
#
# Wizard name: Install
# Wizard source: setup.jar
# Created on: Mon May 09 16:35:04 MST 2005
# Created by: InstallShield Options File Generator
#
# This file can be used to create an options file (i.e., response file) for the
# wizard "Install". 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.
#
#####

#####
#
# License Agreement State
#
# The initial state of the License Agreement panel. Legal values are:
#
# 0 - Nothing will be selected
# 1 - "I accept the terms of the license agreement." will be selected
# 2 - "I do not accept the terms of the license agreement." will be
# selected
#
# For example, to configure the panel to initially display "I accept the
# terms of the license agreement.", use
#
# -W license.selection=2
#
# -W license.selection=1
#
#####
#
# User Input Field - Win_Novell_select
#
# This field only has meaning if you are running the installer on Windows.
# You must choose whether it should install the Windows CLI or the Novell
# CLI.
#
# If you choose Novell, then make sure that the installLocation (see below)
# is on a Novell drive.
#
# 1 - Novell Netware
# 2 - Windows
#
# -W win_user_input.Win_Novell_select="2"
#
#####

```

```

IBM System Storage DS command-line Interface 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="C:\Program Files\IBM\dscli"
#####
#
# User Input Field - os400_java_location
#
#
# This is only used on i5/OS iSeries system installations.
#
#
### -W os400_java_location_panel.os400_java_location="<value>"
#
#####
#
# User Input Field - novell_location_mapped
#
# This is only used on Novell installations.
#
#
### -W novell_user_input.novell_location_mapped="<value>"
#
#####
#
# User Input Field - novel_java_location
#
# This is only used on Novell installations.
#
#
### -W novell_user_input.novel_java_location="<value>"
#
#####

```

5. Save the modifications to the response file.
6. Type the following command for all systems (including UNIX and Linux) except the i5/OS iSeries system at the command prompt `setup<platform>`. `<exe | bin | sh> -silent -options options.txt` and press the **Enter** key on your keyboard to start the installation process in unattended (silent) mode.
7. Type the following command if you are installing the DS CLI onto an i5/OS iSeries system `setup<platform>`. `<exe | bin | sh> -os400 -silent -options options.txt` and press the **Enter** key on your keyboard to start the installation process in unattended (silent) mode.

You can verify that the command-line interface has installed correctly by reviewing the CLI.CFG file in the lib/ subdirectory of the installation directory.

Installing the DS CLI using the console mode

Complete this task to install the DS CLI application on your system using the console mode. The console mode is primarily used for installations on a Linux operating system or on a UNIX operating system without an X display. Users can run the installer from a command prompt on a Windows operating system.

Consider the following before you perform the installation of the DS CLI application:

- You cannot install the DS CLI on a computer that contains the ESS CLI.
- You cannot install the DS CLI on a Windows 64-bit operating system.

- You must have a version of Java 1.4.1 or higher that is installed on your system in a standard directory. The DS CLI installer checks the standard directories to determine if a version of Java 1.4.1 or higher exists on your system. If this version is not found in the standard directories, the installation fails. If the installation fails because the correct version of Java is not found, see *“Correcting the JVM-not-found error”* on page 125.

The DS CLI installation CD contains the Java 1.4.2 version. You can install this version from the CD and continue with the installation of the DS CLI application.

- If the DS CLI application was installed on your client or host machine in the past, ensure that you end any active sessions of the DS CLI before you run the DS CLI installation CD.
- If you are installing onto a Novell system, you are directed to provide the following information:
 - The directory where your Windows drive is mapped
 - The directory where the JAVA HOME is located
- If you are installing the DS CLI onto an OS/400 system, ensure that the machine that you are installing from is network-attached to the OS/400 machine.
- During the installation of the DS CLI application onto an OS/400 system, you must provide the following information:
 - The name of the OS/400 machine to which you are installing the DS CLI application.
 - The user name and password that is used to access the designated OS/400 machine.

The console mode is primarily used for installations on a Linux operating system or on a UNIX operating system that does not have an X display. Users can run the installer from a command prompt on a Windows operating system. The console mode installer displays its various screens as text.

Before you can use the DS CLI, some operating systems require that you restart your system after the installation is complete. Or, you might be required to open a new command prompt window to start a DS CLI session.

Perform the following steps to install the DS CLI using the console mode:

1. Insert the DS CLI installation CD into the CD drive. If the Installshield starts, click the **Exit** selection to exit.
2. Open a command prompt and locate the `setupwin32console.exe` or the `setuplinux.bin` file on the DS CLI CD.
3. Type the following command on the command line (unless you are installing onto an OS/400 system): `setup platform<.exe | .bin | .sh> -console`. The Welcome screen is displayed.

For example, for Windows, type: `setupwin32console.exe` or, for Linux, type: `setuplinux.bin -console`

For an installation onto an OS/400 system from a Windows operating system, type: `setupwin32console.exe -os400`

Notes:

- a. When you issue the `setupwin32console.exe` command on a Windows operating system, you do not need to include the **-console** parameter as part of your command.

- b. Initially the DS CLI installer checks your standard directories for the correct version of Java. If the correct version of Java is not found, you receive the following message:

A suitable JVM could not be found. Please run the program again using the option `-is:javahome <JAVA HOME DIR>`

If you receive this message, see “Correcting the JVM-not-found error” on page 125.

```
Initializing InstallShield Wizard...
Searching for Java(tm) Virtual Machine...
.
Searching for Java 1.4.1 by IBM Corporation
Verifying Java 1.4.1 by IBM Corporation
.....
-----
Welcome to the InstallShield Wizard for IBM System Storage DS Command-Line
Interface (CLI)
The InstallShield Wizard installs IBM System Storage DS Command-Line Interface
on your computer.

To continue, choose Next.

DS Command-Line Interface
IBM Corporation

Users are encouraged to read the User's Guide PDF and the README file before
installing. Both files are located in the cliReadmes/ directory of this
CD-ROM.

Press 1 for Next, 3 to Cancel or 4 to Redisplay [1]
```

4. Press 1 for Next. The License agreement screen is displayed. You can also press 3 to Cancel, or 4 to Redisplay.

```
Please read the following license agreement carefully.

Use of the IBM System Storage DS Command-Line Interface (CLI) is governed by the
IBM Agreement for Licensed Internal Code, a copy of which has been provided
with your DS Machine.

(C) Copyright 2004, 2005 International Business Machines
Corporation All rights reserved.

Please choose from the following options:

[ ] 1 - I accept the terms of the license agreement.
[ ] 2 - I do not accept the terms of the license agreement.

To select an item enter its number, or 0 when you are finished: [0] 1

[X] 1 - I accept the terms of the license agreement.
[ ] 2 - I do not accept the terms of the license agreement.

To select an item enter its number, or 0 when you are finished: [0] 1

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]
```

5. Press 1 to accept the terms of the license agreement. A confirmation prompt is returned with an "X" next to the acceptance line and a 0 in the finished line.

- a. Press Enter to notify the installer that you are finished with this section. The next action line of instructions is displayed.
- b. Press 1 for Next. One of the following results occurs depending on the system that you are using to install the DS CLI:
 - The Windows User's screen is displayed if you are installing on a Windows operating system or are pointing to a Novell system. Continue with the next step.
 - The Install Location screen is displayed if you are installing on a Linux operating system. Continue with Step 10 on page 117.
 You can also press 3 to Cancel, or 4 to Redisplay.

```

Please select the appropriate target system:

[X] 1 - Windows
[ ] 2 - Novell Netware

To select an item enter its number, or 0 when you are finished: [0]

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

```

6. If you want the Windows operating system, press 1. If you want the Novell operating system, press 2. If you choose 1 or 2, a confirmation prompt is returned with an "X" next to the system that you choose and with a 0 in the finished line.
 - a. Press Enter to notify the installer that you are finished with this section.
 - b. Press 1 to continue. The Install Location screen is displayed. This screen can be different depending whether you choose Windows, Novell, or are installing for an OS/400.
 - Continue to Step 7 if you choose Windows.
 - Continue to Step 8 on page 117 if you choose Novell.
 - Continue to Step 9 on page 117 if you are installing for an OS/400.
 You can also press 2 to go to the previous screen, 3 to Cancel, or 4 to Redisplay.
7. (Windows and Linux only) Observe the default directory path that is displayed in the Install Location screen and take the following action (example: the Install Location screen at the end of this step):
 - a. Press 1 to continue if you are satisfied with the default directory.
 - b. Change the directory path where you want to install the CLI (do not use blank spaces in the path) if you are not satisfied, and press 1 to continue.
 - c. Go to Step 10 on page 117 after pressing 1.

The Install Location screen is displayed after pressing 1.

 You can also press 2 to go to the previous screen, 3 to Cancel, or 4 to Redisplay.

```

IBM System Storage DS Command-Line Interface Install Location

Please specify a directory or press Enter to accept the default directory.

Directory Name: [C:\Program Files\ibm\dsccli]or [opt/ibm/dsccli]

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

```

8. (Novell only) Observe the default directory paths that are displayed in the Installation Location screen for Novell and perform the following action (example: the Install Location screen at the end of this step):

Notes:

- a. SYS represents where the Windows drive is mapped in a Novell system. However, this could also be SYS:\DS\CLI or a volume other than SYS:
 - b. J: (could be any drive letter)
- a. Press 1 to continue if you are satisfied with the default values.
 - b. Enter the applicable value if you are not satisfied with the default, and press 1 to continue.
 - c. Go to Step 10 after you press 1.

The Install Location screen is displayed after you press 1.

You can also press 2 to go to the previous screen, 3 to Cancel, or 4 to Redisplay.

```
IBM System Storage DS Command-Line Interface Install Location

Please indicate the Novell location (volume:\directory) where this window J:
drive is mapped:

Enter a value: [SYS:]

Please indicate the JAVA HOME location on Novell: (JVM 1.4.1 or above is
required)

Enter a value: [SYS:\JAVA]

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]
```

9. (Windows and OS/400 only) Record the default directory path that is displayed in the Install Location screen for OS/400 and perform the following action (example: the Install Location screen shown at the end of this step):
 - a. Press 1 to continue if you are satisfied with the default value.
 - b. Enter the applicable value if you are not satisfied with the default, and press 1 to continue.
 - c. Go to Step 10 after you press 1.

The Install Location screen is displayed after you press 1.

You can also press 2 to go to the previous screen, 3 to Cancel, or 4 to Redisplay.

```
Please indicate the directory where Java is installed on the OS/400

Enter a value: [/QOpenSys]

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]
```

10. Verify that the installation directory shown on the Install Location screen is the one you want the DS CLI application installed in. Press 1 to continue. The Installation Progress screen is displayed.

Note: Depending on the current configuration of your system, the total size of the installation can vary.

You can also press 2 to go to the previous screen, 3 to Cancel, or 4 to Redisplay.

IBM System Storage DS Command-Line Interface will be installed in the following location:

C:\Program Files\ibm\dscli

for a total size:

16.6 MB

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

11. Observe the installation progress and wait until the DS CLI installation is completed before pressing 1. Press 1 after the installation completes (there is no reminder to press 1 when the installation is completed). The Install Completed screen is displayed after you press 1.

Installing command-line Interface. Please wait...

Installing IBM System Storage DS Command-Line Interface. Please wait...

|-----|-----|-----|-----|
0% 25% 50% 75% 100%
|||||

|-----|-----|-----|-----|
0% 25% 50% 75% 100%
|||||

Creating uninstaller...

12. Press 1 to continue if you are satisfied with the installation. The Readme screen is displayed (not shown).
You can also press 3 to Cancel, or 4 to Redisplay.

The InstallShield Wizard has successfully installed IBM System Storage DS Command-Line Interface. Choose Next to continue the wizard.

Press 1 for Next, 3 to Cancel or 4 to Redisplay [1]

13. Press Enter to read the readme information or press 3 to finish the installation process. If you press Enter, you must press q to quit the readme file and then press 3 to finish the installation process. The Readme screen contains specific information about the DS CLI for your operating system.

Note: If you are installing on a Windows operating system, you might have to restart your system. If this is the case, the following screen is displayed and you must complete the next step.

To complete the installation you must restart your computer.

1 - Yes, restart my computer.

2 - No, I will restart my computer at a later time.

To select an item enter its number, or 0 when you are finished:

14. Select the restart computer item that you want and press 3 to finish the installation process. Use of the DS CLI application becomes effective with a reboot of your system.

You can verify that the command-line interface has installed correctly by reviewing the CLI.CFG file in the lib subdirectory of the install directory that you defined in Step 6.

Installing the DS CLI using the graphical mode

Complete this task to install the DS CLI application on your system using the graphical installation mode.

Users of Windows, Novell, UNIX, and System i systems can install the DS CLI using the graphical mode. Consider the following before you perform the installation of the DS CLI application:

- You cannot install the DS CLI on a computer that contains the ESS CLI.
- You cannot install the DS CLI on a Windows 64-bit operating system.
- You must have a version of Java 1.4.1 or higher that is installed on your system in a standard directory. The DS CLI installer checks the standard directories to determine if a version of Java 1.4.1 or higher exists on your system. If this version is not found in the standard directories, the installation fails.

The DS CLI installation CD contains the Java 1.4.2 version. You can install this version from the CD and continue with the installation of the DS CLI application.

- If the DS CLI application was installed on your client or host machine in the past, ensure that you end any active sessions of the DS CLI before you run the DS CLI installation CD.
- If you are installing onto a Novell system, you are directed to provide the following information:
 - The directory where your Windows drive is mapped
 - The JAVA HOME directory where the version of Java that you want to use is located.
- System i and i5/OS installation considerations:
 - The latest Java group program temporary fixes (PTF)
 - i5/OS 5722-SS1 option 34 - Digital certificate manager
 - Licensed product 5722-AC3 option *base - Crypto Access Provider 128 bit
 - Licensed product 5722-DG1option *base - IBM HTTP Server for iSeries
 - Licensed product 5722-JV1 options 6 - Java Developer Kit 1.4
 - The latest cumulative PTF package installed on the i5/OS
 - The machine that you are installing from must be connected to the i5/OS through an IP connection.

You can install the DS CLI application using the graphical mode with the help of an installation wizard. Before you can use the DS CLI, some operating systems require that you restart your system after the installation is complete. Or, you might be required to open a new command prompt window to start a DS CLI session.

Note: After you install the new version of DS CLI, your old DS CLI sessions might be unusable.

Perform the following steps to install the DS CLI application using the graphical mode:

1. Start the setup file that is appropriate for your operating system.

If you are ultimately installing the DS CLI onto a System i system, add **-os400** to your command that includes the setup file. For example, `setupwin32.exe -os400`

Initially the DS CLI installer checks your standard directories for the correct version of Java. If the correct version of Java is not found, you receive the following message:

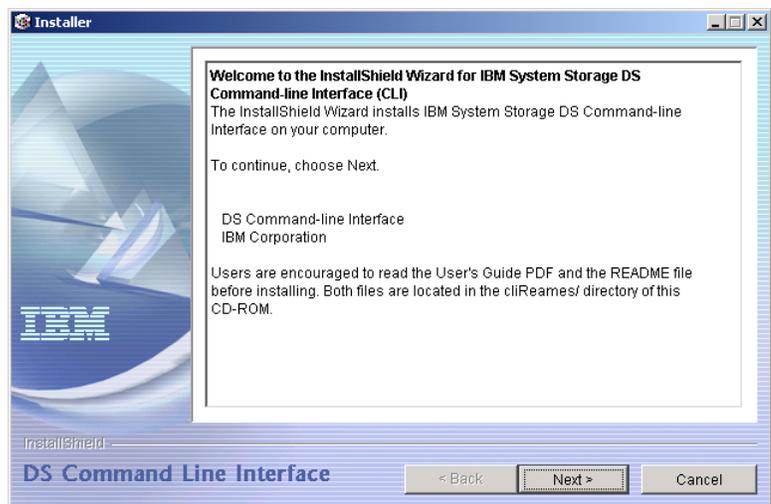
A suitable JVM could not be found. Please run the program again using the option `-is:javahome <JAVA HOME DIR>`

If you receive this message, see “Correcting the JVM-not-found error” on page 125.

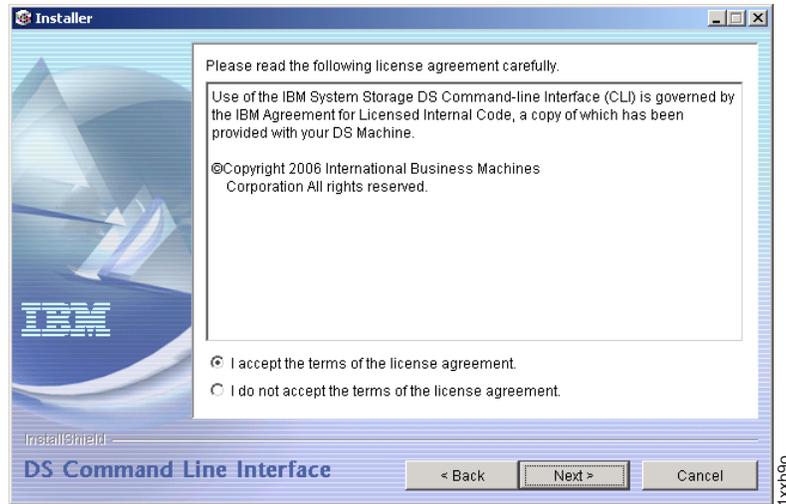
Go to Step 2 if you are installing for a System i system.

Go to Step 3 for all installations other than System i.

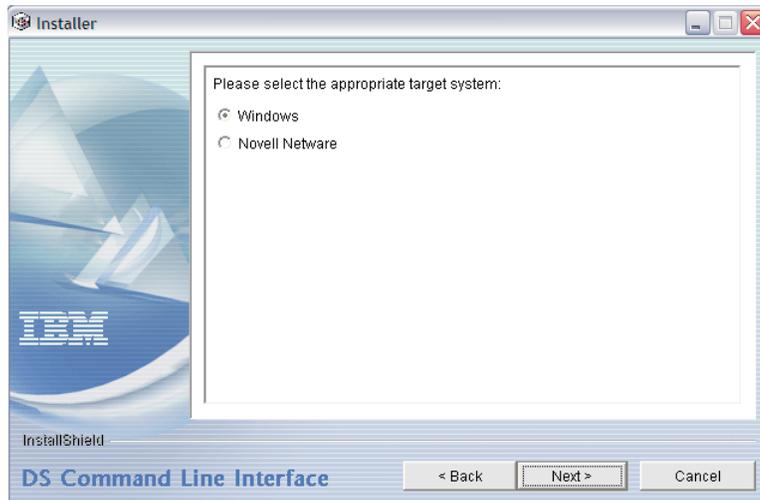
2. The installation wizard displays a sign-on panel that directs you to provide the name of the System i system to which you are installing the DS CLI application. You must provide the IP address or the DNS name of the i5/OS server. In addition, you must also supply the user name and password that is used to access the designated System i system. Click **OK**.
3. The Welcome window opens. Click **Next** to continue or **Cancel** to exit the installation.



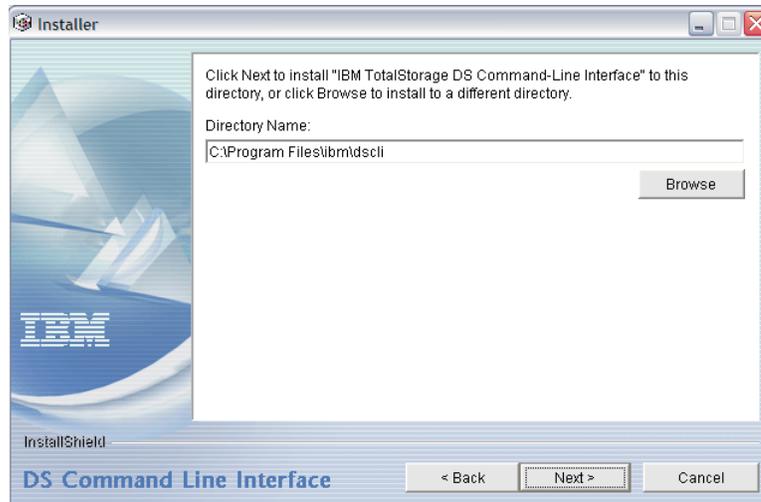
4. The License Agreement window opens. Select “**I accept the terms of this license agreement**” to continue. Click “**I do not accept the terms of this license agreement**” or **Cancel** to exit the installation.



5. The Target System window opens. Select the appropriate target system where you want the DS CLI installed, and then click **Next** to continue or **Cancel** to exit the installation.



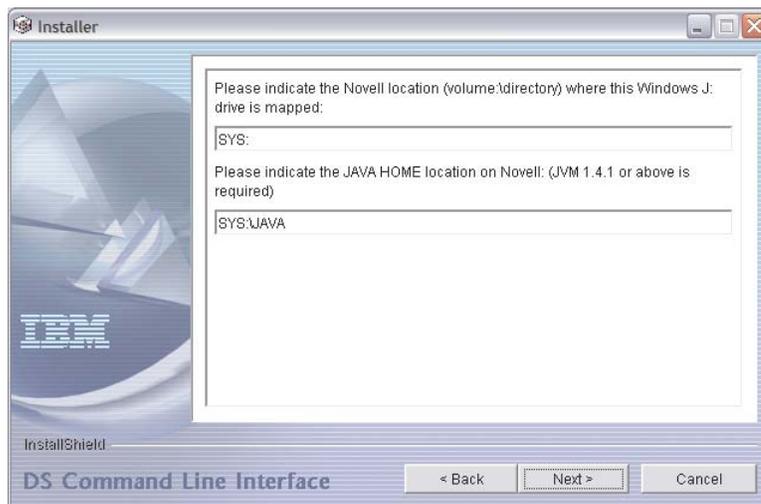
6. The Installation Location window (Windows has been selected as your operating system) opens. If you selected Novell, go to Step 7 on page 122. If you are installing on a Linux system, go to Step 9 on page 123. If you are satisfied with the Windows default directory, click **Next**. If not, type the directory path where you want to install the DS CLI and click **Next**. Do not use blank spaces in the installation path. Click **Cancel** if you want to exit the installation.



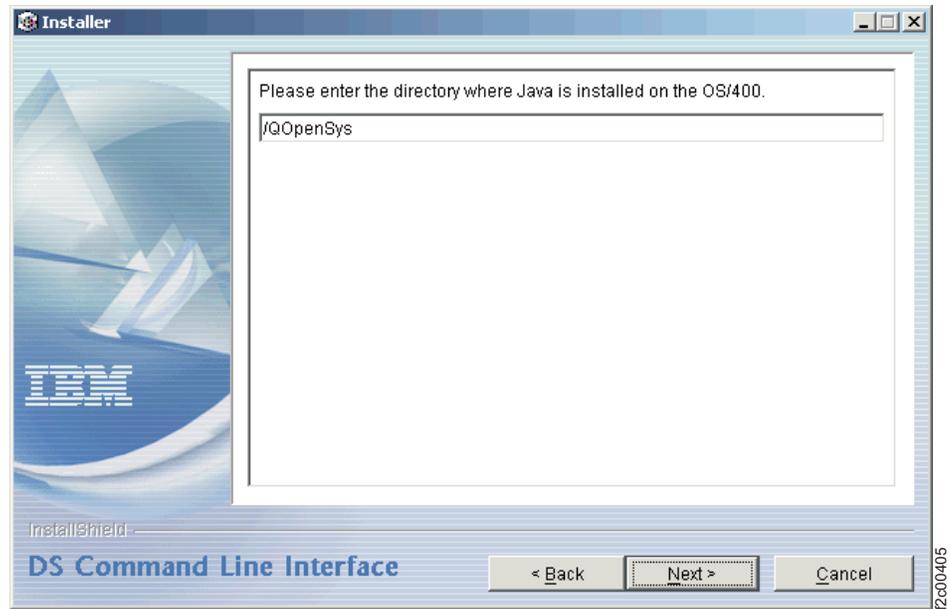
|
|
|

Note: If you are installing onto a System i system, a window that asks for the directory where Java is installed on the i5/OS is displayed when you click **Next**. Go to Step 8.

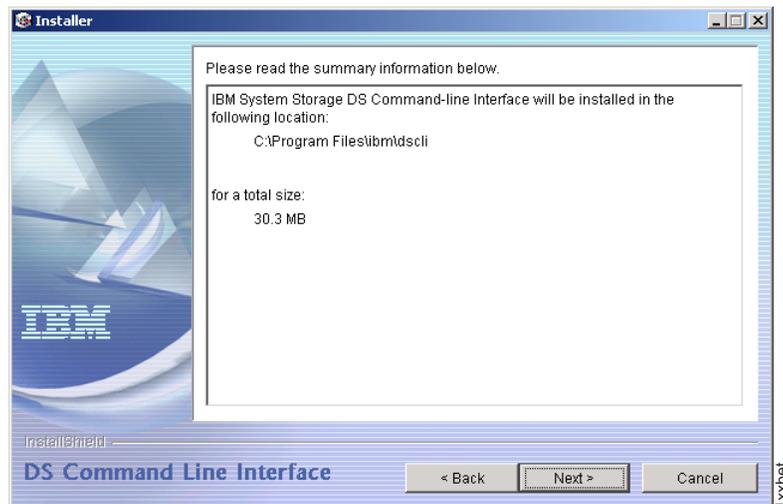
7. The Novell Location window opens. Complete the information for where the Windows drive is installed and where the JAVA HOME directory that contains the version of Java you want to use is located. Click **Cancel** if you want to exit the installation.



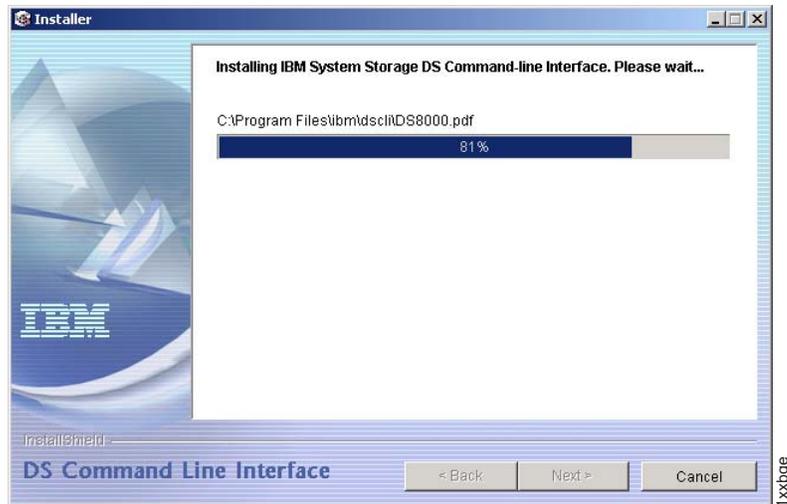
8. The Java directory window for the System i system opens. Insert the directory information. Click **Next** to continue or **Cancel** to exit the installation.



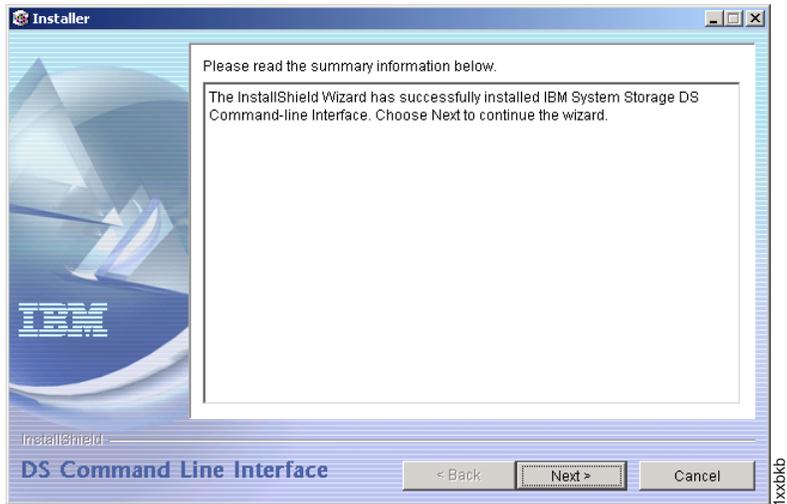
9. The Installation Preview window opens. This window provides the location where the command-line interface will be installed and specifies how much space it will occupy on your drive. Click **Next** to continue or **Cancel** to exit the installation. You can change the installation directory by clicking **Back**.



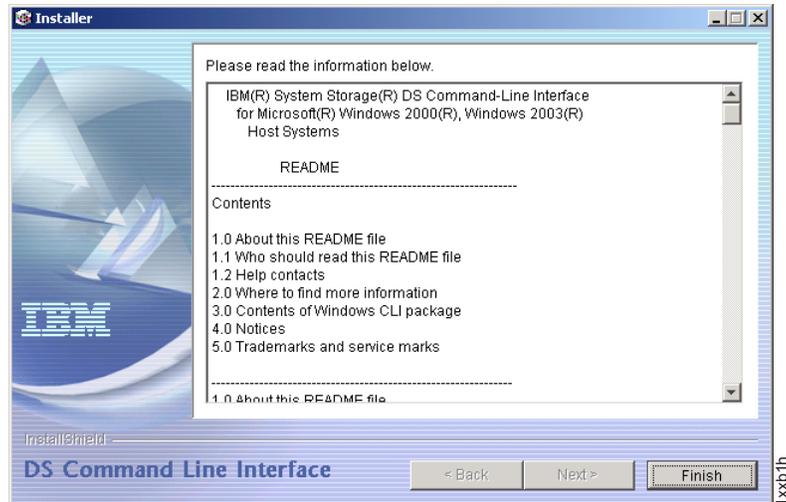
10. The Installation Progress window opens. This window provides the progress of the command-line interface installation. Click **Next** to continue or **Cancel** to exit the installation.



11. The Installation Summary window opens. This window provides the installation summary information. Click **Next** to continue or **Cancel** to exit the installation.



12. The Installation Complete window opens. A reminder is provided in text to read the Readme file. Click **Finish** to complete the installation.



Notes:

1. You can verify that the command-line interface has installed correctly by reviewing the CLI.CFG file in the lib/ directory.
2. The DS CLI application is installed in the following two places in i5/OS:
 - IFS directory *IBM/DS_CLI*. This directory contains the profiles, exe files, java jars, readme files, and so forth.
 - The QDSCLI library. This library contains executable code.
3. Before you can invoke the DS CLI application from the i5/OS, you must add the QDSCLI library to the i5/OS library list.

Correcting the JVM-not-found error

Complete this task to correct the JVM-not-found error.

The JVM-not-found error occurs if the DS CLI installer cannot find the correct version of Java in the standard directories of your system. You must have Java 1.4.1 or higher on your system for the DS CLI application to work.

Notes:

1. This error might also occur if you are installing the DS CLI on an AIX system. The LIBPATH environment variable can interfere with the installation of the DS CLI and results in the display of the JVM-not-found error. It is recommended that you disable the LIBPATH environment variable before you install the DS CLI. After the installation of the DS CLI, enable the LIBPATH environment variable so that it can be used with other applications.
2. If you have already installed the DS StorageManager, you have installed Java 1.4.2 on your system. Because it was not installed in one of your standard directories, the DS CLI installer did not find it during its initial check.

If Java 1.4.1 is not found during the initial check, the following message is displayed:

```
A suitable JVM could not be found. Please
run the program again using the option -is:javahome <JAVA HOME DIR>
```

Perform either step 1 or step 2 of the following steps to correct the JVM not found error (see Step 4 for AIX):

1. (Optional) Run the installer again, and provide the path to the JVM. The following examples represent directories where you can find the correct version of Java:

For a Windows system, specify the following:

```
setupwin32 -is:javahome  
C:\Program Files\IBM\dsniserver\java\jre
```

For a Linux system, specify the following:

```
setuplinux.bin -is:javahome  
/opt/ibm-Java-whatever/java/bin
```

Note: If you use this argument, the installer attempts to use whatever JVM that you specify, even if it is an unsupported version.

Go directly to step 4 after this step completes.

2. (Optional) Open the DS CLI installation CD and install the Java 1.4.2 version into one of your standard directories.
3. Continue with the installation of the DS CLI application.
4. (AIX) Run the following commands to sequentially disable the LIBPATH environment variable, install the DS CLI, and restore the LIBPATH environment variable:

```
export LIBSAVE=$LIBPATH  
unset LIBPATH  
setupaix.bin -is:javahome /opt/ibm-Java-whatever/java/bin  
export LIBPATH=$LIBSAVE  
unset LIBSAVE
```

Proceed to the next step after you disable the LIBPATH environment variable.

Installing the DS CLI on your OpenVMS system

Use the following steps to install the DS CLI application in your OpenVMS environment.

Before you install the DS CLI application, 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.

1. Log on to your host system as a user with SYSLOCK, SYSNAM, SYSPRV, (or a system UIC), TMPMBX, and CMKRNL privileges.
2. Insert the DS CLI product CD-ROM into the CD drive.
3. Mount the CD drive. For example, for an IDE CD device DQA0, type the following command: **MOUNT /NOASSIST /OVERRIDE=IDENTIFICATION /MEDIA_FORMAT=CDROM DQA0**: A message similar to the following is displayed.

```
%MOUNT-I-WRITELOCK,volume is write locked  
%MOUNT-I-CDROM_ISO, : (1 of 1) , mounted on VMS1$DQA0:
```

For a SCSI CD device DKA *nnn*, type the following command: **MOUNT /NOASSIST /OVERRIDE=IDENTIFICATION /MEDIA_FORMAT=CDROM DKA_{nnn}**: where *nnn* represents the number that is assigned by the OpenVMS system to your CD device.

4. Type the following command: **DIRECTORY/FULL DQA0: [000000] IBM-AXPVMS-DSCLI-* .PCSI** and press **Enter** to access the command-line interface installation

package in the root directory of the CD. Output similar to the following is displayed.

```
Directory DQA0:[000000]
IBM-AXPVMS-DSCLI-V0500-01F96-1.PCSI;1  File ID: (4,7,0)
Size: 55.79MB/55.80MB Owner: [0,0]
Created: 9-MAR-2005 04:07:22.25
Revised: 9-MAR-2005 04:09:43.98 (1)
Expires: None specified
Backup: No Backup record
Effective: None specified
Recording: None specified
Accessed: None specified
Attributes: None specified
Modified: None specified
Linkcount: 1
File organization: Sequential
Shelved state: Online
Caching attribute: Writethrough
File attributes: Allocation: 114282, Extend: 0, Global buffer count: 0
Version limit: 0, Backups disabled
Record format: Undefined, maximum 0 bytes, longest 0 bytes
Record attributes: None
RMS attributes: None
Journaling enabled: None
File protection: System:RWED, Owner:RWED, Group:RWED, World:RWED
Access Cntrl List: None
Client attributes: None

Total of 1 file, 55.79MB/55.80MB
```

5. Type the following command and press Enter to extract the command-line interface for OpenVMS release notes: **PRODUCT EXTRACT RELEASE_NOTES DSCLI /SOURCE=DQA:[000000] /FILE=filespec**

Note: If you do not use the **/FILE** qualifier, the release notes are written to the DEFAULT.PCSI\$RELEASE_NOTES file in your current default directory. Read the release notes before continuing to ensure that you are equipped with the information that you need to complete your installation successfully.

6. Type the following command and press Enter to invoke the command-line interface installation process: **PRODUCT INSTALL DSCLI /SOURCE=DQA0:[000000] /DESTINATION=devicename:[directoryname] /RECOVERY_MODE**. A message similar to the following is displayed.

```
The following product has been selected:
IBM AXPVMS DSCLI V5.0-1F96 Layered Product
Do you want to continue? [YES]
```

The **/DESTINATION** qualifier can be omitted from the command. However, IBM recommends that you use the optional **/RECOVERY_MODE** qualifier. For a detailed description of all **PRODUCT INSTALL** command qualifiers and parameters, see the HP OpenVMS System Management Utilities Reference Manual or the OpenVMS online help.

7. Press Enter to continue the installation process. The following message is displayed.

```
Configuration phase starting ...

You will be asked to choose options, if any, for each selected product and for
any products that may be installed to satisfy software dependency requirements.

IBM AXPVMS DSCLI V5.0-1F96: IBM DS Command Line Interface (DS CLI)

(C) Copyright IBM Corp. 1999, 2005. All Rights Reserved.

International Business Machines Corporation (IBM)

No PAK

Do you want the defaults for all options? []
```

8. Type No and press Enter to review the installation options.

Note: You can type Yes if you have decided to install with the default options.

The following CLI license message is displayed after typing Yes or No.

```
License Agreement IBM System Storage DS Command-line Interface

Copyright 2005 International Business Machines Corporation
All rights reserved.

Use of the IBM System Storage DS Command-line Interface (DS CLI) is
governed by the IBM Agreement for Licensed Internal Code, a copy
of which has been provided with your DS Machine.

Answer "Yes" to install the DS CLI code. By installing this code,
you are certifying that you have read and accept the IBM DS CLI
License agreement.

Answer "No" to terminate and exit the IBM DS CLI installation.

Do you want to continue? [YES]
```

9. Press Enter to accept the DS CLI license agreement. The following Java license message is displayed.

License Agreement HP Java Run-Time Environment for OpenVMS

The IBM DS CLI requires the Java 1.4.2 Java Runtime Environment (JRE). Installing the DS CLI program automatically installs the required JRE. The installed JRE is the intellectual property of and licensed by Hewlett-Packard Company.

You can view the license agreement on the World Wide Web at:

http://h18012.www1.hp.com/java/download/ovms/1.4.2/rte1.4.2_license.html

Answer "Yes" to install the Java code. By installing this code, you are certifying that you have read and accepted the HP Java License agreement.

Answer "No" to terminate and exit the IBM DS CLI installation.

Do you want to continue? [YES]

10. Press Enter to accept the Java license agreement. The command-line interface requires Java 1.4.2 on an OpenVMS Alpha host system. If you answer No, the installation process automatically ends and exits. The following library update warning message is displayed.

WARNING: By default the system Help and Message libraries will be updated.

The IBM DS CLI program provides local Help and Help Message library files. By default, the CLI installation integrates these local libraries into the OpenVMS system Help and Help Message libraries.

To prevent the system libraries from being modified chose to review the installation options and answer "No" when prompted to update the libraries.

Do you want to continue? [YES]

11. Press Enter to continue the installation process. The following configuration question is displayed.

Would you like the local IBM DS CLI Help and Help Message libraries to be integrated into the OpenVMS system libraries?

If you answer "Yes", the following OpenVMS libraries will be updated:

SYS\$COMMON:[SYSHLP]HELPLIB.HLB
SYS\$COMMON:[SYSHLP]MSGHLP\$LIBRARY.MSGHLP\$DATA

If you answer "No", OpenVMS system libraries will not be modified.

In every case, local libraries are available under:

IBMDSCLI\$HELP:IBMDSCLI_OVR.HLB
IBMDSCLI\$HELP:IBMDSCLI_MESSAGES.MSGHLP\$DATA

Modify the system Help and Help Message libraries? [YES]

12. Press Enter to confirm the library update option (or type No and press Enter to deny the library update option). The following confirmation message is displayed:

Do you want to review the options? [N0]

13. Press Enter to confirm and accept all selections. The following installation message with completion status is displayed.

```
Execution phase starting ...

The following product will be installed to destination:
  IBM AXPVMS DSCLI V5.0-1F96
  DISK$V732_ALPHA:[VMS$COMMON.]

Portion done: 0%...10%...20%...30%...60%...70%...80%...90%...100%

The following product has been installed:
  IBM AXPVMS DSCLI V5.0-1F96          Layered Product
```

14. Review the Installation Verification Procedure (IVP) report, which is similar to the following output, and check for possible errors.

```
%PCSI-I-IVPEXECUTE, executing test procedure for IBM AXPVMS DSCLI
V5.0-1F96 ...
dscli -ver
IBM DSCLI Version: 5.0.1.96
%PCSI-I-IVPSUCCESS, test procedure completed successfully
```

15. Ensure that the installation completes. When the Polycenter Software Installation utility finishes the command-line interface installation process, a message similar to the following is displayed.

```
IBM AXPVMS DSCLI V5.0-1F96: IBM DS Command Line Interface (DS CLI)

Insert the following lines in SYS$MANAGER:SYSTARTUP_VMS.COM:
  @PCSI$DESTINATION:[IBMDSCLI.MGR]IBMDSCLI$STARTUP.COM
Insert the following lines in SYS$MANAGER:SYSHUTDOWN.COM:
  @IBMDSCLI$MANAGER:IBMDSCLI$SHUTDOWN.COM

Users of this product require the following lines in their login
command procedure:
  @IBMDSCLI$MANAGER:IBMDSCLI$LOGIN.COM

Release notes for IBM DS CLI available in IBMDSCLI$HELP
```

16. Unmount the CD drive and remove the CD-ROM.

The command-line interface provides program startup, login, and shutdown procedures in the [destinationdir.IBMDCLI.MGR] directory. The installation process runs the startup and login procedures immediately before invoking the IVP procedure. But for persistent setup, you must integrate the startup, login, and shutdown procedures.

Using the DS CLI on an OpenVMS system

To use the DS CLI on an OpenVMS system, you must integrate the required IBMDSCLI\$STARTUP.COM and IBMDSCLI\$LOGIN.COM procedures. You can also integrate the optional IBMDSCLI\$SHUTDOWN.COM procedure. This integration is accomplished when you use the OpenVMS persistent setup procedure.

You must add the required IBMDSCLI\$STARTUP.COM procedure to your system startup processes. The IBMDSCLI\$STARTUP.COM procedure defines the logical names that are required for the command-line interface in your system logical name table and installs some images with enhanced privileges. This procedure is intended to be invoked during the system startup.

You also must add the required IBMDSCLI\$LOGIN.COM procedure to your system login processes. The IBMDSCLI\$LOGIN.COM procedure sets up the JRE that is required by the command-line interface and defines the DSCLI command as a foreign DCL command symbol. It is intended to be invoked during the system-wide SYS\$MANAGER:SYLOGIN.COM or user-specific SYS\$LOGIN:LOGIN.COM procedure.

In order to use the DS CLI interface from start to finish in your OpenVMS system, you can add the optional IBMDSCLI\$SHUTDOWN.COM procedure to your system shutdown processes. The IBMDSCLI\$SHUTDOWN.COM procedure performs the removal operations for privileged images and undefines system-wide logical names that are associated with the DS CLI. It is intended to be invoked during the system shutdown process.

1. Integrate the IBMDSCLI\$STARTUP.COM procedure into your system startup by adding the following line to the SYS\$MANAGER:SYSTARTUP_VMS.COM script: `$ @destinationdev:[destinationdir.IBMDSCLI.MGR] IBMDSCLI$STARTUP`

destinationdev

Specifies the name of the device that contains the command-line interface installation directory

destinationdir

Specifies the name of the directory where you just installed the command-line interface.

Note: You can alternatively add the IBMDSCLI\$STARTUP.COM procedure to the SYSMAN startup database.

2. Integrate the IBMDSCLI\$LOGIN.COM procedure into the system-wide or user-specific login by adding the following line:

`$@IBMDSCLI$MANAGER:IBMDSCLI$LOGIN.COM`

Note: Run the IBMDSCLI\$LOGIN.COM procedure only after you have successfully run the IBMDSCLI\$STARTUP.COM procedure.

3. Integrate the IBMDSCLI\$SHUTDOWN.COM procedure by adding the following line to the SYS\$MANAGER:SYSHUTDOWN.COM script:

`$@IBMDSCLI$MANAGER:IBMDSCLI$SHUTDOWN`

Note: This step is optional. However, processing this step allows your system to make the full use of the DS CLI application.

Completing DS CLI postinstallation

Complete these tasks to set up the DS CLI application so that you can use the DS CLI to configure your DS6000.

You must install the DS CLI before you complete these postinstallation tasks.

After you install the DS CLI application, how do you log into the application? After the initial login, what are some of the first tasks that you must accomplish before you can get the full benefit of the DS CLI application? The instructions in this section answer these questions, as well as, describe how to configure your volumes and how to manage potential problems.

Before you can use the DS CLI application with the i5/OS (particularly if you are using external load source), you must set up the initial configuration of your DS6000 models. After the initial configuration, you can do a D-mode IPL and begin using the DS CLI application directly from the i5/OS. In the meantime, you can follow the list below as a guide for your initial configuration.

1. Initiate the DS CLI to begin using it in either single-shot, script, or interactive command mode.
2. Set up your required user accounts and passwords.
3. Set your DS CLI default configuration settings.
4. Activate your licensed functions. This includes obtaining your feature activation codes and applying the feature activation codes to your storage unit.
5. Enable the remote support and call home functions. You must provide customer contact and network information to enable these functions.
6. Register for the My Support service.
7. Configure new fixed block storage. Use the DS CLI to create and modify fixed block extent pools, arrays, ranks, volumes, and volume groups. You can also configure host ports and connections.

Using the DS CLI application

You must ensure that you have installed the DS Storage Manager using the full-management console installation and that you have configured your domain. Without this domain configuration (which is a one-time process), you cannot use the DS CLI. After you install the DS CLI, three command modes are available to you.

You must log into DS CLI to use the command modes. There are three command modes for the DS CLI:

- Single-shot
- Interactive
- Script

Note: You must not start more than 100 DS CLI sessions simultaneously. Starting more than 100 DS CLI sessions simultaneously can result in connection problems.

Logging into the DS CLI application

You must log into the DS CLI application to use any of the command modes.

You must ensure that you are in the directory where you installed the DS CLI application. The following list provides a reminder of the supported operating systems default directories where the DS CLI is installed if the directory designation is not changed:

AIX /opt/ibm/dscli

HPUX /opt/ibm/dscli

Linux /opt/ibm/dscli

Sun Solaris
/opt/ibm/dscli

Windows
C:\Program Files\IBM\dscli

HP Tru64
/opt/ibm/dscli

iSeries
/ibm/dscli

Novell Netware
SYS:\dscli

When you log into the DS CLI application (type `dscli` at the command prompt), you must provide the following information:

HMC1
Specify the primary management console.

User Name
Specify the name of the user account. The default account for the first login is **admin**.

Password
Specify the user password. The default password for the admin account is **admin**. However, this password is only good for your first login.

Note: Because the password for the admin account expires after you log in for the first time, you must change the password before you can perform any other DS CLI command function. Use the **chuser** command to change your password.

The first time that you log in to the DS CLI, you can specify this information using either of the following two methods:

- Ensure you are in the directory where you installed the DS CLI application and type the `dscli` command at the command prompt. Supply all the log in information with the command. For example: `dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t`.
Use this command when you use the single-shot mode for the first time and when the DS CLI application is not active on your system. In addition, when you use the single-shot mode, you must include the command that you want to process. For example, if you want to process the **lssi** command, if you have not activated the DS CLI application, and if you are using the single-shot mode, type: `dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t lssi`.
- When you log into the DS CLI application (from the directory where you installed the DS CLI application) by typing `dscli`, you are prompted to supply the information for HMC1, user name, and password.

Notes:

1. Entering a DS CLI command at the dscli command prompt requires that you continue entering all the parameters and values until the command is complete. This can result in an automatic line wrap if your command has many parameters and values.
2. You cannot force a line break or wrap by hitting the Enter key and then entering the rest of the command on a second line. The DS CLI application interprets the use of the Enter key as an end to the function and begins to process whatever is contained on the line, ignoring the second line.
3. The DS CLI command examples that are provided in this guide are often shown with line wraps that would not occur during your input. These examples are displayed for clarity and other formatting considerations.

Using the DS CLI single-shot command mode

Use the DS CLI single-shot command mode if you want to issue an occasional command but do not want to keep a history of the commands that you have issued.

You must supply the login information and issue the command that you want to process at the same time. Perform the following steps to use the single-shot mode:

1. Use the following command format to enter a DS CLI session (Windows operating system):

```
dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t  
lssi -s -fullid -hdr off
```

Here is an example of this same command in i5/OS without the report delimiters:

```
DSCLI SCRIPT(*NONE) HMC1('9.5.17.156') USER(admin) PASSWORD(itso4a11)  
DACL(lssi)
```

This command demonstrates the use of the **lssi** command with the **-s** parameter. Use this command to view the storage image IDs for your storage complex. The storage image ID consists of the manufacture name (IBM), the machine type (1750), and the serial number.

Notes:

- a. The command example uses the **-fullid** DS CLI command flag. The **-fullid** command flag generates fully qualified IDs, which include the storage image ID, for every ID that is displayed in the command output.
 - b. The command example also uses the **-hdr off** command flag which turns off the header that is generally associated with the report that is generated from the **lssi** command.
 - c. Almost every DS CLI command requires the use of the storage image ID. If you add your target storage ID into your profile file under the **-devid** designation, you are not required to provide the storage image ID when you issue a command. However, if you do provide the **-dev** (*storage_image_ID*) parameter, the value that you type takes priority over the value contained in your profile file.
2. Wait for the command to process. The following type of report is generated that lists the storage image IDs that are associated with the storage complex.

IBM.1750-68FA111

IBM.1750-68FA111

IBM.1750-68FA120

Using the DS CLI script command mode

Use the DS CLI script command mode if you want to issue a sequence of DS CLI commands. Administrators can use this mode to create automated processes; for example, establishing remote mirror and copy relationships for volume pairs.

Consider the following when using the DS CLI script command mode:

- The DS CLI script can contain only DS CLI commands. Use of shell commands results in a process failure.
- You can add comments to the scripts. Comments must be prefixed by the number sign (#); for example, # This script contains PPRC Path establish procedures.

Note: It is not the intent of this instruction to tell you how to write a script. An example script is displayed for your use as a guide.

You can issue the DS CLI script from the command prompt at the same time that you provide your login information.

1. Type the script name at the command prompt using the following format:
`dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd password -script ~/bin/mkpprcpairs`

Note: If you are using i5/OS and have already logged on to the DS CLI application, you invoke the script mode using the following format:
`DSCLI SCRIPT('/myscript') USER(admin) OUTPUT('/outfile')`

2. Wait for the script to process and provide a report regarding the success or failure of the process.

The following is an example of a script that could be used to establish remote mirror and copy relationships for volume pairs.

```
mkpprc -dev IBM.1750-1303561 -remotedev IBM.1750-7504491 -type mmir 1000-103F:2300-233F
mkpprc -dev IBM.1750-1303561 -remotedev IBM.1750-7504491 -type gcp 1100-113F:2340-237F
mkpprc -dev IBM.1750-1303561 -remotedev IBM.1750-7504491 -type mmir 1800-187F:2800-287F
mkpprc -dev IBM.1750-1303561 -remotedev IBM.1750-7504491 -type gcp 1200-127F:2500-257F
mkpprc -dev IBM.1750-1303561 -remotedev IBM.1750-7504491 -type mmir 1040-1054:2700-2714
mkpprc -dev IBM.1750-1303561 -remotedev IBM.1750-7504491 -type gcp 1055-107F:2400-242A
mkpprc -dev IBM.1750-1303561 -remotedev IBM.1750-7504491 -type mmir 1140-117F:2600-263F
```

Using the DS CLI interactive command mode

Use the DS CLI interactive command mode when you have multiple transactions to process that cannot be incorporated into a script. The interactive command mode provides a history function that makes repeating or checking prior command usage easy to do.

In addition to being able to enter DS CLI commands at the DS CLI command prompt, a history function provides a view of the last DS CLI commands that you have used. It also allows you to repeat any of the last commands more quickly than having to type out the entire command. The example at the end of this process shows how the history function works.

Perform the following steps to use the DS CLI in the interactive command mode:

1. Log on to the DS CLI application at the directory where it is installed.

Note: If you make a mistake and type the wrong user name or password, do not try to correct the error within the current session. Exit the DS CLI session that you are in and log in to a new DS CLI session.

2. Provide the information that is requested by the information prompts. The information prompts might not appear if you have provided this information in your profile file. The command prompt switches to a **dscli** command prompt.
3. Begin using the DS CLI commands and parameters. You are not required to begin each command with **dscli** because this prefix is provided by the **dscli** command prompt.

Note: Issue the **setoutput** command if you know that you will be using interactive mode for a lengthy period of time. The **setoutput** command allows you to set or display command output format options. You control how the reports that are generated by the list and the show commands are displayed on your computer. For example, you can specify that the reports be displayed in one of the following formats:

- **delim** = displays output in a table format and sets the column delimiter to a single character
- **xml** = displays output in XML format
- **stanza** = displays output in stanza (horizontal table) format

See the **setoutput** command for more details.

To use the DS CLI history function that is associated with the interactive command mode, perform the following steps:

1. Issue an exclamation mark (!) to display CLI commands that you have used in the current session. For example: **dscli>!** results in a list of commands such as the following:

```
[4] lsarraysite -dev IBM.1750-1300771
[3] lsarray -dev IBM.1750-1300771
[2] lsextpool -dev IBM.1750-1300771
[1] lsextpool -dev IBM.1750-1300771
```

2. Issue **dscli>!1** to retry the last command. Or, issue **dscli>!3** to retry command [3].

DS CLI command help

You can obtain online help for each CLI command. To obtain the help, type the word *help* and the command name at the **dscli** command prompt. There are additional parameters that you can use with the help command that designate the type of help information that you can receive.

The **help** command contains the following parameters that influence the type of help information that you can receive:

Command	Description
help	Displays a list of all the DS CLI commands that are available for use.
help -s	Displays a list of commands with brief descriptions.

Command	Description
help -l	Displays a list of commands with their associated syntax.
command_name -h command_name -help command_name -? help command_name	Displays the reference page (man page) for the command name.
help -s command_name	Displays the brief description for the command name.
help -l command_name	Displays the usage statement for the command name.

Note: You cannot use the **-s** and **-l** parameters with the following help command flags: **-h**, **-help**, and **-?**.

Examples

The following examples represent the type of information that is displayed when you use various parameters with the help command. Each of these examples start at the dscli command prompt.

Note: Much of the information that is associated with the **help** command is displayed in list format. You can include the page (**-p on**) and row (**-r number**) controls; for example, `dscli>help -p on -r20`. This command pauses your page listing after 20 entries and prompts you to **press any key to continue**.

Example 1

```
dscli>help
```

This line of input provides the entire list of DS CLI commands. Only the command names are displayed. No other details are provided.

Example 2

```
dscli>help -s
```

This line of input provides the entire list of DS CLI commands and a short description for each command. The short description explains what each command accomplishes.

Example 3

```
dscli>help -l
```

This line of input provides the entire list of DS CLI commands and the syntax for each command. In addition, you see all the help syntax and formatting parameters, which can make your reading of the help difficult. For example, you can see something similar to the following for each command:

```
lsxtpool [ { -help|-h|-? } ] [ { -l (long)|-s (short) } ]
[-fmt default|xml|delim|stanza] [-p on|off] [-delim char]
[-hdr on|off] [-bnr on|off] [-r #] [-v on|off] [-fullid]
[-hmc1 SMC1] [-hmc2 SMC2] [-user username] [-passwd password]
-dev storage_image_ID [-stgtype fb|ckd] [-rankgrp 0|1]
[Extent_Pool_ID ... | -]
```

The following line is the actual syntax for the **lsextpool** command: `-dev storage_image_ID [-stgtype fb|ckd] [-rankgrp 0|1] [Extent_Pool_ID ... | -]`

Example 4

```
dscli>lssi -h
dscli>lssi -help
dscli>lssi -?
dscli>help lssi
```

Any of the previous lines of input generates the entire help page that is associated with the **lssi** command. This is the same information that is found in the *IBM System Storage DS6000 Command-Line Interface User's Guide* or in the Information Center for the **lssi** command.

Example 5

```
dscli>help -s lssi
```

This line of input provides the short description that is associated with the designated command. For instance, you might see the following:

```
lssi
```

The **lssi** command displays a list of storage images in a storage complex. You can use this command to look at the status of each storage image in the list. The storage image worldwide node name (WWNN) is displayed when this command is used. You must use the storage image WWNN when using the **lsavailprcport** and **mkpprcpath** commands.

Example 6

```
dscli>help -l lssi
```

This line of input displays the entire reference page found in both the *IBM System Storage DS6000 Command-Line Interface User's Guide* and in the Information Center for the **lssi** command.

Obtaining the serial (storage image ID) number using the DS CLI

Almost every DS CLI command requires the use of the storage image ID. If you add your target storage image ID into your profile file under the *-devid* designation, you are not required to provide the storage image ID when you issue each command.

Use the **lssi** or **lssu** DS CLI commands to list the storage image IDs that are associated with your storage complex. It is recommended that you record the target storage image ID in your profile file. This can save you input time when you have to process many transactions that cannot be part of a script.

To obtain a list of the storage image ID numbers, perform the following steps. The storage image ID consists of the manufacture name (IBM), the machine type (1750), and the serial number.

1. Log into the DS CLI application in interactive command mode (if you are not already logged in).
2. Type the following command format at the dscli command prompt to obtain the storage image IDs. Substitute **lssi** or **lssu** for the command name variable in the example.

```
dscli>commandname -s -fullid -hdr off
```

Note: The result of using the parameters and command flags as shown in the example is the same for either command. However, each command can produce different information when you use their associated parameters. See the **Issi** or **Issu** commands for more details.

3. Wait for the command to process. The following type of report is generated, which lists the storage image IDs that are associated with the storage complex.

IBM.1750-75FA111
IBM.1750-75FA111
IBM.1750-75FA120

Activating your machine and feature licenses using the DS CLI

Use this scenario to activate your license activation codes. These codes must be activated before any configuration can be applied to your DS6000 network.

The following licenses can be activated depending on your purchase:

- Operating environment license for each storage unit that you own. (This license must be activated.)
- Copy Services, which can consist of the following features:
 - Point-in-time
 - Remote mirror and copy

There are multiple codes associated with these features. To obtain the information that you need to activate these licenses and features in your storage unit, go to the IBM Disk Storage Feature Activation (DSFA) Web site at:

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

Download your codes onto a diskette in XML format. You can then import the codes from the XML file when you process the DS CLI **applykey** command.

Note: In most situations, the DSFA application can locate your order confirmation code (OCC) when you enter the DS6000 (1750) serial number and signature. However, if the OCC is not attached to the 1750 record, you must assign it to the 1750 record in the DSFA application. In this situation, you must have the OCC (which you can find on the License Function Authorization document).

Use the **applykey** DS CLI command to activate the licenses for your storage unit. Use the **lskey** DS CLI command to verify which type of licensed features are activated for your storage unit.

1. Log into the DS CLI application in interactive command mode (if you are not already logged in).
2. Issue the DS CLI **applykey** command at the dscli command prompt as follows (this example presumes that your XML file is named "keys" and it resides on a diskette in your A: drive):

```
dscli> applykey -file a:\keys.xml
IBM.1750-68FA120
```
3. Press Enter. When the process has completed, the following message is displayed:

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI
Version 5.0.0.0 DS: IBM.1750-68FA120
Licensed Machine Code key xxxx, key xxxx successfully applied.
```

4. Verify that the keys have been activated for your storage unit by issuing the DS CLI **lskey** command as follows: `lskey IBM.1750-68FA120`
5. Press Enter and the following type of report is displayed:

```
Sun Apr 09 02:23:49 PST 2006 IBM DS CLI
```

Activation Key	Authorization Level (TB)	Scope
Operating Environment (OEL)	45	All
Remote mirror and copy (RMC)	25.1	All
Remote mirror for z/OS (RMZ)	25.1	CKD
Point in time copy (PTC)	On	All
Parallel access volumes (PAV)	On	CKD

Setting up user accounts using the DS CLI

This scenario describes how to set up a user account. You must have administrator authority to enable this function.

The admin account is set up automatically at the time of installation. It is accessed using the user name **admin** and the default password **admin**. This password is temporary and expires after its initial use. You must change the password before you can use any of the other functions. There are 7 groups the administrator can assign to a user. A user can be assigned to more than one user group. The groups and the associated functions allowed by the assignment are as follows:

admin

All users that you assign to the administrator user group allows access to all storage management console server service methods and all storage image resources.

op_volume

The volume operator user group allows access to service methods and resources that relate to logical volumes, hosts, host ports, logical subsystems, logical volumes, and volume groups, excluding security methods. In addition, this user group inherits all authority of the monitor user group.

op_storage

The storage operator user group allows access to physical configuration service methods and resources, including storage complex, storage image, array, rank, and extent pool objects. This user group inherits all the authority of the `op_copy_services` and monitor user groups, excluding security methods.

op_copy_services

The copy services operator user group allows access to all Copy Services service methods and resources, excluding security methods. In addition, this user group inherits all authority of the monitor user group.

service

The service user group includes monitor authority, plus access to all management console server service methods and resources, such as performing code loads and retrieving problem logs.

monitor

The monitor user group allows access to list and show commands. It provides access to all read-only, nonsecurity management console server service methods and resources.

no access

The no access user group does not allow access to any service methods or storage image resources. By default, this user group is assigned to any user account in the security repository that is not associated with any other user group.

In addition to assigning users to one or more user groups, you also must assign a default password to each user. When you notify users of their group assignment and default password, indicate that the default password is only good for the initial log on. Users must change the password at the time of their initial log on. Also, remind all users to record their password in a safe place, because there is no way that the administrator or the application can retrieve a password.

Note: You must change the default password for an account, including the **admin** account, to be able to use any CLI command other than the one to change the password. See the `chuser` command for more information.

Use the `mkuser DS` CLI command to create new user accounts with specific roles (user group or groups) and an initial password. If you assign multiple roles to an account, ensure that you separate the different roles by using a comma for example, `op_volume, op_storage`. See the `mkuser` command description for more details.

1. Log into the DS CLI application in interactive command mode.
2. Issue the following command from the `dscli` command prompt to assign a user to an account with a default password: `dscli>mkuser -pw AB9cdefg -group service,op_copy_services testuser`
3. Press Enter and observe the processing result. A successful process returns the following display:

```
Sun Aug 11 02:23:49 PST 2004 IBM DS  CLI
Version 5.0.0.0 DS: IBM.1750-68FA120
User Name testuser successfully created.
```

Configure support options using the DS CLI

This scenario provides a high-level description of the tasks that you must perform to configure support options using the command-line interface. Support options (for example, e-mail alerts and alert traps that are sent by an SNMP agent) enable IBM support personnel to quickly assist you with problem determination and, with your consent, remotely perform certain maintenance procedures.

Note: The Call Home function also has a modem call home feature. The DS CLI cannot be used to configure this feature. You must use the DS Storage Manager to configure the modem call home feature.

Before you begin, you must have the command-line interface prompt, and you must be connected to a storage unit that is used for open systems host system storage.

This scenario first creates the contact information and then establishes the SMTP port and IP address where a message is sent in the event of a problem. After these

two tasks are completed, you can activate the Call Home function and verify that the Call Home function works according to your specifications. You can also set up your storage unit to receive alert traps.

1. Set up the contact information.

The contact information that you supply is sent to IBM so that an IBM service representative can contact you. The type of information that you must supply consists of the following but is not limited to this list:

- Company name
- Company address
- Country
- State or province
- Primary contact phone number
- Primary e-mail address (required if you are using the Call Home e-mail notification feature).

See “Setting up contact information using the DS CLI” for details.

2. Set up the e-mail notification process.

Using e-mail, the Call Home function can notify IBM when a problem occurs on your storage unit. You must specify the SMTP port and IP address where the message is sent in the event of a problem.

See “Setting up Call Home (SMTP) notifications using the DS CLI” on page 144 for details.

3. Activate and verify the Call Home function.

After you provide your contact information and specify the SMTP port and IP address where the message is sent in the event of a problem, you can activate and test Call Home.

See “Activating the Call Home function using the DS CLI” on page 159 for details.

4. Set up the SNMP alert notifications process.

One of the support options you can configure for your storage unit involves the notification of Simple Network Management Protocol (SNMP) alert traps. SNMP notifications are alerts and notifications of thresholds that have been exceeded and exception events that have occurred.

See “Setting up SNMP notifications using the DS CLI” on page 146 for details.

5. Set up the Service information messages (SIM) notification process.

SIMs are generated by a storage unit for S/390 and zSeries hosts. You can designate the level of error you want to be notified about and how often you want to be notified.

See “Setting up SIM notifications using the DS CLI” on page 146 for details.

Setting up contact information using the DS CLI

Complete this task to establish your contact information in your system. When you use any of the remote support features, the contact information you supply is sent to IBM so that an IBM service representative can contact you.

You **must** supply the following information when you are setting up for remote support:

- Company name
- Ship state or province
- Ship location
- Ship country

- Contact primary phone number
- Machine identification (consists of the manufacture, machine type, and serial number)
- Primary e-mail address (required if you are using the Call Home e-mail notification feature).

It is also recommended that you supply the following additional types of information when setting up for remote support:

- Company address
- Machine location (the physical location of the machine)
- Machine address (where the machine is physically located)
- Ship phone (the telephone number of the person to contact when parts are being shipped)
- Ship to city
- Ship postal code
- Contact name (the name of the system administrator who can be contacted by IBM service personnel)
- Contact alternate phone number

Use the **setcontactinfo** command to provide the contact information that is needed for the remote support function.

Perform the following steps to establish your contact information in the system. The example commands that are displayed in this task are shown in two formats. The first format provides the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **setcontactinfo** command to supply the contact information. Enter the **setcontactinfo** command at the dscli command prompt with the following parameters and variables:

```
dscli>setcontactinfo -companyname company_name -shipstate [state, province]
                    -shiploc building_location -shipcountry country
                    -contactpriphone primary_phone_number
                    -contactemail primary_email_address storage_image_ID
```

Example

```
dscli>setcontactinfo -companyname Examples -shipstate AZ
                    -shiploc bld1300 -shipcountry US -contactpriphone 520-799-8001
                    -contactemail example@example.com IBM.1750-68FA120
```

Notes:

- a. The phone number is displayed with hyphens that separate the various parts.
 - b. If you decide to use the company address, ensure that you use double quotation marks around your input. For example:
-companyaddr "9210 S. High Rd Tucson Az 85744"
2. Press Enter. The following message is displayed if the process is successful:


```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.1750-68FA120
```

The contact information settings were successfully modified
 3. Issue the **showcontactinfo** command to verify that the contact information is active in your system. Enter the **showcontactinfo** command at the dscli prompt with the following parameter:

```
dscli>showcontactinfo storage_image_ID
```

Example

```
dscli>showcontactinfo IBM.1750-68FA120
```

A report is displayed showing the contact information that is established and recognized by your system.

Setting up Call Home (SMTP) notifications using the DS CLI

Complete this task to set up the Call Home function to contact IBM support using e-mail when a problem occurs on your storage unit. The **setsmtp** command is used to set up the e-mail feature of the Call Home function.

Note: The Call Home function also has a modem call home feature. The DS CLI cannot be used to configure this feature. You must use the DS Storage Manager to configure the modem call home feature.

Consider the following requirements before you use the **setsmtp** command to set up e-mail notifications:

- You must know your SMTP port and IP address, because these are required to activate the e-mail notifications. This information specifies where a message is sent in the event of a problem.
- Your system must recognize all the required contact information before you can issue the **setsmtp** command. This includes the following information:
 - Company name
 - Ship state
 - Ship location
 - Ship country
 - Primary contact phone number
 - Primary e-mail address (required if you are using the call home e-mail notification feature).

Perform the following steps to activate the e-mail notification support option feature. The example commands that are displayed in this task are shown in two formats. The first format provides the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **showplex** command to obtain a report that provides your SMTP port and IP address. Enter the **showplex** command at the dscli command prompt with the following parameter:

```
dscli>showplex storage_image_ID
```

Example

```
dscli>showplex IBM.1750-68FA120
```

2. View the report and make a copy of it. The important thing is that you have a record of the SMTP and SNMP IP addresses and port numbers that are used in setting up the Call Home function.
3. Issue the **showcontactinfo** command. This command displays a report that allows you to ensure that the required contact information, including the primary e-mail address, has been set in the system. Enter the **showcontactinfo** command at the dscli command prompt with the following parameter:

```
dscli>showcontactinfo storage_image_ID
```

Example

```
dscli>showcontactinfo IBM.1750-68FA120
```

4. Review the report and ensure that all the required contact information has been set in the system, including the primary e-mail address. If any of the required information is missing, issue the **setcontactinfo** command before you proceed with this task.

5. Issue the **setsmtp** command to specify where a message is sent in the event of a problem. Enter the **setsmtp** command at the dscli command prompt with the following parameters and variables:

```
dscli>setsmtp -server smtp_server_IP:port_number -emailnotify [on, off]
storage_image_ID
```

Example

```
dscli>setsmtp -server 9.11.242.12:25 -emailnotify on
IBM.1750-68FA120
```

6. Press Enter. The following message is displayed if the process is successful:

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.1750-68FA120
```

```
The SMTP settings were successfully modified
```

Setting up your account information using the DS CLI

Complete this task to set up your account information which is needed when using the Call Home function. You must use the **setplex** command to associate a unique name with the 1750, which in turn allows you to open a secure shell connection to a storage unit for remote support.

Before you can activate and test the Call Home function, ensure that the following tasks have been completed:

- You have supplied the following required contact information using the **setcontactinfo** command:
 - Company name
 - Ship state
 - Ship location
 - Ship country
 - Primary contact phone number
 - Primary e-mail address (required if you are using the call home e-mail notification feature).
- You have supplied the location where e-mail notifications are sent by specifying the SMTP port and IP address through the **setsmtp** command.

To set up your account information, provide the following information:

- Account number
- The secure shell connection enabled
- Storage image ID (manufacture, machine type, and serial number) unless you have stored this information in your profile file

Note: The example in the following steps presumes that the storage image ID information has not been stored in the profile file.

Perform the following steps to set up your account information. The example commands that are displayed in this task are shown in two formats. The first format provides the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **setplex** command to establish your account information in the system. Enter the **setplex** command at the dscli command prompt with the following parameters and variables:

```
dscli>setplex -acct account -allowrssh enable storage_image_ID
```

Example

```
dscli>setplex -acct 101010 -allowrssh enable IBM.1750-68FA120
```

2. Press Enter. The following message is displayed if the process is successful:

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.1750-68FA120
```

```
Storage Plex successfully modified
```

3. Issue the **showplex** command to verify that your system recognizes the account information. Enter the **showplex** command at the dscli command prompt with the following parameter:

```
dscli>showplex storage_image_ID
```

Example

```
dscli>showplex IBM.1750-68FA120
```

A report is displayed that details the properties of your 1750 storage complex. Detailed properties include names, descriptions, and notification settings for the storage complex.

Setting up SNMP notifications using the DS CLI

One of the support options you can configure for your storage unit involves the notification of Simple Network Management Protocol (SNMP) alert traps. SNMP notifications are alerts and notifications of thresholds that have been exceeded and exception events that have occurred. This task shows you how to set up your storage unit to receive these notifications as they occur.

You need to know your SNMP port and IP address as these are required when activating the notifications feature. This information specifies where a message is sent in the event of a problem.

To use the **setsnmp** command to activate the SNMP alert notification support option feature, perform the following steps:

1. Log in to the DS CLI application in interactive command mode, if you are not already logged in.
2. Specify where a message is sent in the event of a problem by issuing the **setsnmp** command at the dscli command prompt as follows:

```
dscli>setsnmp -server 9.11.242.12:161 IBM.1750-68FA120
```

3. Press Enter and a successful process displays the following message:

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.1750-68FA120
```

```
The SNMP settings were successfully modified
```

Setting up SIM notifications using the DS CLI

Service information messages (SIMs) are generated by a storage unit for S/390 and zSeries hosts. You can designate the level of error you want to be notified about and how often you want to be notified.

It is recommended that you determine your levels of notification before using the **setsim** command. You need to decide what type of message should generate a

message. You have the following choices including a level of severity, severity designations are the same for each type of message. You can also designate how often you want to be notified.

SIM message types consist of the following:

- DASD service information messages
- Media service information messages
- Storage unit messages

SIM severity designations consist of the following choices:

- acute - An irrecoverable error with possible loss of data.
- moderate - A system path is not operational and performance might be degraded.
- serious - An irrecoverable error or a data check with loss of access to data.
- service - A recoverable error, equipment checks, or data checks. You can defer repair.
- none

SIM notifications for each message type can be in the range of 0-5 time.

To set up your SIM notification choices using the `setsim` command, perform the following steps:

1. Log in to the DS CLI application in interactive command mode, if you are not already logged in.
2. Specify your notification choices using the following example format for the `setsim` command. At the `dscli` command prompt type the following:

```
dscli>setsim -dasdlevel serious -dasdnotify 3 -medialevel moderate  
-medianotify 2 -sulevel service -sunotify 1 IBM.1750-68FA120
```

3. Press Enter and a successful process displays the following message:

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.1750-68FA120
```

```
The SIM settings were successfully modified
```

Creating a default configuration setup with a profile file

You can create default settings for the command-line interface by defining one or more profiles on the system. For example, you can specify the storage management console (SMC) for the session, the output format for list commands, the number of rows per page in the command-line output, and the banner that is included with the command-line output.

If a user enters a value with a command that is different from a value in the profile, the command overrides the profile.

You have several options for using profile files:

- You can modify the default profile. The default profile, `dscli.profile`, is installed in the profile directory with the software. For example, `c:\Program Files\IBM\DSCLI\profile\dscli.profile` is the directory path for the Windows platform and `/opt/ibm/dscli/profile/dscli.profile` is the directory path for UNIX and Linux platforms.
- You can create a personal default profile by making a copy of the system default profile as `<user_home>/dscli/profile/dscli.profile`.

- You can create a profile for the storage unit operations. Save the profile in the user profile directory. For example:

```
<user_home>\dscli\profile\operation_name1
```

```
<user_home>\dscli\profile\operation_name2
```

These profile files can be specified using the DS CLI command parameter **-cfg** *<profile_name>*. If the -cfg profile file is not specified, the user's default profile file is used. If a user's profile file does not exist, the system default profile file is used.

The home directory *<user_home>* is defined by the Java system property named "user.home" The location of your password file is determined by your operating system. The following examples are home directories in different operating systems:

Windows operating system

For a Windows operating system, the property value defaults to the environment variable %USERPROFILE%. As a result, your personal profile is C:\Documents and Settings\username\dscli\profiles\dscli.profile.

UNIX or Linux operating system

For an UNIX or Linux operating system, the property value defaults to the environment variable \$HOME. As a result, your personal profile is ~/dscli/profile/dscli.profile.

OpenVMS system

For an OpenVMS operating system, the property value defaults to the logical name SYS\$LOGIN. As a result, your personal profile is [.dscli.profile]dscli.profile.

Note: The values of the Java system properties can be redefined by JRE options. If you are having problems, check to see if you have an environment setting like the following on your local system:

```
_JAVA_OPTIONS=-Duser.home=...
```

When you install the command-line interface software, the default profile is installed in the profile directory with the software. The file name is dscli.profile. For example. c:\Program Files\IBM\DSCLI\profile\dscli.profile.

Table 11 provides the list of profile variables that you can use to create the profile.

Table 11. Profile variables

Variable	Description
banner: on off	Enables or disables the banner that appears before the command output. This variable is equivalent to the command option -bnr . The command option -bnr overrides this default value.
delim	Specifies a delimiter character for the format: delim variable. The default character is a comma. This variable is equivalent to the command option -delim . The command option -delim overrides this default value.
devid	Specifies the storage image ID that is the target for the command. This value is equivalent to the command option -dev . The command option -dev overrides this default value.

Table 11. Profile variables (continued)

Variable	Description
format	<p>Specifies the output format for list commands.</p> <p>Specify one of the following formats:</p> <ul style="list-style-type: none"> • default: Specifies default output. • xml: Specifies XML format. • delim: Specifies columnar format. Columns are delimited with the character that you must specify with the delim variable. • stanza: Specifies a horizontal table. <p>This variable is equivalent to command option -fmt. The command option -fmt overrides this default value.</p>
fullid	<p>Specifies that IDs display in fully qualified format, which includes the storage image ID.</p>
header: on off	<p>Enables or disables the headers that display with the columns of data in the list commands. This variable is equivalent to the command option -hdr. The command option -hdr overrides this default value.</p>
hmc1	<p>Specifies the primary Storage Manager IP address. This variable is equivalent to the command option -hmc1. The command option -hmc1 overrides this default value.</p>
hmc2	<p>Specifies the secondary Storage Manager IP address. This variable is equivalent to the command option -hmc2. The command option -hmc2 overrides this default value.</p>
locale	<p>Specifies the language for the output on the local computer.</p> <ul style="list-style-type: none"> • ar: Arabic • be: Byelorussian • bg: Bulgarian • ca: Catalan • cs: Czech • da: Danish • de: German • el: Greek • en: English • es: Spanish • et: Estonian • fi: Finnish • fr: French • gu: Gujarati • hi: Hindi • hr: Croatian • hu: Hungarian • in: Indonesian • is: Icelandic • it: Italian • iw: Hebrew • ja: Japanese • kk: Kazakh • kn: Kannada • ko: Korean • lt: Lithuanian • lv: Latvian (Lettish) • mk: Macedonian • mr: Marathi • ms: Malay

Table 11. Profile variables (continued)

Variable	Description
locale, <i>continued</i>	<ul style="list-style-type: none"> • nl: Dutch • no: Norwegian • pa: Punjabi • pl: Polish • pt: Portuguese • ro: Romanian • ru: Russian • sa: Sanskrit • sh: Serbo-Croatian • sk: Slovak • sl: Slovenian • sq: Albanian • sr: Serbian • sv: Swedish • ta: Tamil • te: Telugu • th: Thai • tr: Turkish • uk: Ukrainian • vi: Vietnamese • zh: Chinese
paging: on off	Controls the display of output. If paging is enabled, a limited number of lines of output displays when a command is issued. The lines do not scroll. You must set the number of lines per page with the rows variable. This variable is equivalent to command option -p . The command option -p overrides this default value.
timeout	Set timeout value of client/server synchronous communication. The unit of the value is second. The default value is 420 seconds. You can set this timeout if the processing of a command ends by timeout due to network or client/server performance issue. Note: The command timeout value can be longer than this value because one command can consist of multiple client/server requests.
timeout.connection	Set timeout value to establish client or server connection. The unit of this value is seconds. The timeout value must be greater than zero. System default socket timeout value is used if the value is set to zero. The default value is 20 seconds. Notes: <ol style="list-style-type: none"> 1. If the DS CLI returns a connection error, check for the following conditions: <ul style="list-style-type: none"> • Is there a secure physical connection between the client and server? • Is the default timeout value too short to establish a connection? 2. Setting a connection timeout value too short can cause unexpected connection problems.

Table 11. Profile variables (continued)

Variable	Description
remotedevid	Specifies the remote storage image ID. This variable is equivalent to the command option -remotedev . The command option -remotedev overrides this default value.
rows	Specifies the number of rows per page of output if the paging variable is enabled. This variable is equivalent to command option -r . The command option -r overrides this default value.
verbose: on off	Enables or disables verbose output. This variable is equivalent to the command option -v . The command option -v overrides this default value.

Example

```
#
# DS CLI Profile
#

#
# Management Console/Node IP Address(es)
# hmc1 and hmc2 are equivalent to -hmc1 and -hmc2 command options.
#hmc1: 127.0.0.1
#hmc2: 127.0.0.1

#
# Default target Storage Image ID
# "devid" and "remotedevid" are equivalent to
# "-dev storage_image_ID" and "-remotedev storage_image_ID" command options,
# respectively.
#devid: IBM.1750-AZ12341
#remotedevid: IBM.1750-AZ12341

#
# locale
# Default locale is based on user environment.
#locale: en

# Timeout value of client/server synchronous communication in second.
# DSCLI command timeout value may be longer than client/server communication
# timeout value since multiple requests may be made by one DSCLI command
# The number of the requests made to server depends on DSCLI commands.
# The default timeout value is 420 seconds.
#timeout: 900

# Socket connection timeout value in seconds.
# The timeout value must be greater than zero.
# System default socket timeout value is used if timeout value is set to zero.
# The default connection timeout value is 20 seconds.
#timeout.connection: 20

# Output settings
#
# ID format of objects:
# on: fully qualified format
# off: short format
fullid: off

# Paging and Rows per page.
# paging enables/disables paging the output per line numbers specified by "rows".
# "paging" is equivalent to "-p on|off" option.
# on : Stop scrolling per output lines defined by "rows".
# off : No paging. (default)
# "rows" is equivalent to "-r #" option.
```

```

paging: off
#rows: 24

# Output format type for ls commands, which can take one of the following values:
# default: Default output
# xml : XML format
# delim : delimit columns using a character specified by "delim"
# stanza : Horizontal table format
# "format" is equivalent to option "-fmt default|xml|delim|stanza".
#format: default

# delimiter character for ls commands.
#delim: |
# Display banner message. "banner" is equivalent to option "-bnr on|off".
# on : Banner messages are displayed. (default)
# off : No Banner messages are displayed.
banner: on

#
# Display table header for ls commands. "header" is equivalent
# to option "-hdr on|off".
# on : Table headers are displayed. (default)
# off : No table headers are displayed.
header: on

#
# Display verbose information. "verbose" is equivalent to option "-v on|off".
# on : Display verbose information.
# off : No verbose information.
verbose: off

# End of Profile

```

OpenVMS system integration

You can adjust your OpenVMS system to obtain greater benefits from the use of the DS CLI application. The hints and tips that are provided in this section show how to obtain these benefits through the optimal integration of the DS CLI into your OpenVMS system.

The following list provides the areas that you might consider for optimizing the use of the DS CLI in your OpenVMS system:

- Command Console LUN (CCL)
- OpenVMS system messages
- Message help
- Java Run Time Environment (JRE)
- Quota recommendations

Enhancing the command console LUN for DS CLI use

The OpenVMS operating system considers a fibre-channel device with LUN ID 0 as Command Console LUN (CCL). These devices do not normally display when you issue the DS CLI **lshostvol** command. However, with adjustments, these devices can be displayed when you issue the **lshostvol** command. The following description provides the information that you need to make this enhancement work on your OpenVMS system.

Fibre-channel CCL devices have the OpenVMS device type GG, which result in OpenVMS device names in the form \$1\$GGAn. In contrast, fibre-channel disk

devices have the OpenVMS device type DG, which result in device names in the form \$1\$DGA*n*. Therefore, LUN 0 devices on OpenVMS are a special device type, different from disk devices.

The DS CLI **lshostvol** command displays the mapping of host device names or volume names to machine type 2105, 2107, and 1750 volume IDs. That implies that all host devices belonging to 2105/2107/1750 volumes are displayed. Therefore, CCL devices \$1\$GGA*n* are included in the **lshostvol** output for multiplatform consistency and to match the output of other DS CLI commands.

However, the inclusion of CCL devices can be confusing for users who expect that the **lshostvol** command displays only the disk devices. You can use the OpenVMS logical name **IBMDSCLI\$SHOW_GG_DEVICES** to modify the DS CLI behavior: If this logical name translates to an expression which evaluates as True in OpenVMS conventions (1, Y, YES, T, or TRUE), then the \$1\$GGA*n* CCL devices are shown in the command output. Otherwise, the \$1\$GGA*n* CCL devices are not shown.

The startup procedure **IBMDSCLI\$STARTUP.COM** defines the logical name **IBMDSCLI\$SHOW_GG_DEVICES** as Y. If you want to suppress \$1\$GGA*n* CCL devices in the **lshostvol** command output, you can redefine the logical name after the startup procedure has been processed.

Enhancing the OpenVMS system messages

When you use the DS CLI, the application provides messages regarding the application processes, status, and errors. You also receive the OpenVMS system messages but they are displayed in a different format. You can make this situation less confusing by making the following adjustments.

The DS CLI messages are presented in an operating-system independent format. In contrast, native OpenVMS programs provide messages using the system message facility as displayed in the following format: **%facility-level-identification, text**.

To ensure that the OpenVMS command **SET MESSAGE** and customer-written tools that scan for such messages work correctly, the DS CLI provides each message using OpenVMS system services in addition to the operating system independent output. After displaying the OpenVMS message, the normal DS CLI message is provided unchanged. This ensures that the DS CLI messages are identical across platforms and that you can work with the DS CLI documentation.

However, these redundant messages can be confusing for users who are not familiar with OpenVMS. You can use the OpenVMS logical name **IBMDSCLI\$OPENVMS_MESSAGES** to modify the DS CLI behavior: If this logical name translates to an expression which evaluates as True in OpenVMS conventions (1, Y, YES, T, or TRUE), then the additional OpenVMS-formatted messages are presented. Otherwise, only the operating system independent DS CLI messages are shown.

The startup procedure **IBMDSCLI\$STARTUP.COM** defines the logical name **IBMDSCLI\$OPENVMS_MESSAGES** as Y. If you want to suppress the OpenVMS-formatted messages, you can redefine the logical name after the startup procedure has been processed.

Enabling OpenVMS to use the DS CLI help

The DS CLI installation process offers the option to add modules to the system help library. If you enable OpenVMS with this option, you can use the DS CLI help.

The DS CLI installation process offers the option to add modules to the system help library `SYS$COMMON:[SYSHLP]HELPLIB.HLB` and the system messages database `SYS$COMMON:[SYSHLP]MSGHLP$LIBRARY.MSGHLP$DATA`. If you choose this option, the module `IBMDSCLI` is added as the top-level key to the help library, and the DS CLI status messages can be accessed using the `HELP/MESSAGE/FACILITY=IBMDSCLI` command. Additionally, the login procedure `IBMDSCLI$MANAGER:IBMDSCLI$LOGIN.COM` activates the message section file `IBMDSCLI$SYSTEM:IBMDSCLI_Messages_Shr.exe` for the current process.

In every case, the installation process provides the following files in the directory which is referred by the logical name `IBMDSCLI$HELP`:

IBMDSCLI_Ovr.hlp

A help library containing one module with the top-level key `IBMDSCLI`. You can add this library to the search list for help libraries in your OpenVMS system by defining appropriate logical names `HLP$LIBRARY`, `HLP$LIBRARY_1`, `HLP$LIBRARY_2`, and so forth.

IBMDSCLI_Messages.msghlp\$data

A message help data file with messages for facility `IBMDSCLI`. You can add this data file to the searchlist for message help files in your OpenVMS system by defining the logical name `MSGHLP$LIBRARY` accordingly.

If you do not want the installation process to modify the OpenVMS system libraries, you can use these OpenVMS default logical names to integrate the DS CLI help information manually.

Java Runtime Environment considerations for DS CLI

The DS CLI login procedure `IBMDSCLI$MANAGER:IBMDSCLI$LOGIN.COM` defines `JAVA$CLASSPATH` in the OpenVMS process logical name table and it overrides any existing Java classpath definition. If you want to use other Java-based software in the same process, you must redefine `JAVA$CLASSPATH` so that it provides the classpath as a `JAVA` command parameter.

The following information provides an overview of how the installation of the DS CLI affects the Java environment of your OpenVMS system.

Because the DS CLI relies on Java Run Time Environment (JRE) V1.4.2, mandatory JRE files are installed in the directory tree that is referenced by the logical name `IBMDSCLI$JRE`. This setup is according to HP guidelines. The login procedure `IBMDSCLI$MANAGER:IBMDSCLI$LOGIN.COM` calls the JRE setup procedure which defines several logical names and DCL symbols for usage by the Fast Virtual Machine.

If your OpenVMS host system uses other software that requires JRE but cannot run with the same JRE version as the DS CLI, users of that software can switch between different Java versions. To use different JRE versions, you must run a command procedure to set up the Java environment definitions for the version that you want to use in the given process (see the OpenVMS Java documentation at:

http://h18012.www1.hp.com/java/documentation/1.4.2/ovms/docs/user_guide.html

The DS CLI application-specific Java classes are bundled in Java Archive (.JAR) files in the directory referenced by logical name **IBMDSCLI\$LIBRARY**. These files must be included in the Java classpath. On OpenVMS, two logical names define the classpath:

CLASSPATH

For UNIX-style names. You can use a string inside single quotation marks that consists of colon-separated path names.

JAVA\$CLASSPATH

For OpenVMS specification syntax. You can specify multiple paths with a comma-separated expression (not enclosed in single quotation marks) as OpenVMS logical name search list. **JAVA\$CLASSPATH** overrides **CLASSPATH**, if **JAVA\$CLASSPATH** is defined.

Because of this override process, you might have to redefine the **JAVA\$CLASSPATH** to provide the class path as a JAVA command parameter. However, this JAVA command parameter is only required if you want to use other Java-based software in the same process.

Quota considerations for DS CLI

The JRE was designed to perform optimally on UNIX systems, where each process is given large quotas by default. On OpenVMS, the default behavior gives each process lower quotas so that many processes can co-exist on a system.

To get the best Java performance on OpenVMS, HP recommends that you set process quotas to match a typical UNIX system. HP also recommends these as minimum quota settings (except where noted). See these recommendations at:

http://h18012.www1.hp.com/java/documentation/1.4.2/ovms/docs/user_guide.html

To check if your current process quotas fulfill the recommendations, you can run the following process: `IBMDSCLI$JRE:[LIB]Java$Check_Environment.com`.

Chapter 6. Enabling support options

This section lists the high-level steps that you must complete in order to enable the remote support, Call Home (SMTP) and SNMP alert options for your IBM System Storage DS6000 series.

You must install the IBM System Storage DS Storage Manager before you can complete these steps. Also, ensure that you meet the VPN connection requirements for remote support.

Each of these steps contains a link to the topic that contains the substeps for each individual task. Follow the link to complete the task and return to this topic to find the link for the next task.

1. Register for My Support. See “Registering for the My Support function” on page 158 for additional steps.
2. Install the modem on the management console. The modem is required for remote support. You previously installed the modem during hardware installation. If necessary, see “Installing a modem on the management console” on page 69 for instructions to install the modem on the management console.
3. Set up the Call Home connection. You initially configured your Call Home connection information during the setup procedures that you completed after installing the DS Storage Manager. You can verify, and if necessary, change your Call Home settings by following the steps in “Setting up Call Home, SNMP, and SIM notifications (real-time only)” on page 100. If you want the modem to be used, check the **Enable Modem Call Home** box on the Configure notifications — Call Home page.

If you are using the DS CLI, you can configure and verify your Call Home settings by following the steps in “Setting up Call Home (SMTP) notifications using the DS CLI” on page 144. You can activate the Call Home function from the DS CLI by following the steps in “Activating the Call Home function using the DS CLI” on page 159, “Setting up contact information using the DS CLI” on page 142 and “Setting up your account information using the DS CLI” on page 145.

Note: You must use the DS Storage Manager to configure the modem call home feature. The DS CLI cannot be used to configure this feature.

4. Set up your SNMP alerts. You initially configured your SNMP alerts during the setup procedures that you completed after installing the DS Storage Manager. You can verify, and if necessary, change your SNMP settings by following the steps in “Setting up Call Home, SNMP, and SIM notifications (real-time only)” on page 100.

If you are using the DS CLI, you can configure and verify your SNMP settings by following the steps in “Setting up SNMP notifications using the DS CLI” on page 146.

5. Set up your zSeries service information messages (SIMs). You initially configured your SIMs during the setup procedures that you completed after you installed the DS Storage Manager. You can verify, and if necessary, change your SIM settings by following the steps in “Setting up Call Home, SNMP, and SIM notifications (real-time only)” on page 100.

If you are using the DS CLI, you can configure your SIM settings by following the steps in “Setting up SIM notifications using the DS CLI” on page 146.

6. Set up the remote support function. You initially configured your contact information during the setup procedures that you completed after you installed the DS Storage Manager. You can verify, and if necessary change your contact information by following the steps in “Defining customer contacts (real-time only)” on page 99.

If you are using the DS CLI, you can configure and verify your contact information by following the steps in “Setting up contact information using the DS CLI” on page 142.

7. Initiate a remote support connection. You can initiate a remote support connection through your local area network or through the modem by following the instructions in “Activating remote support (real-time only)” on page 160. The modem is always used to initiate a VPN connection if there are phone numbers that are configured on the Configure modem remote support page. If you want to initiate a VPN connection through your local area network, you must unconfigure the modem phone numbers. Follow the steps in “Unconfiguring modem phone numbers (real-time only)” on page 160 to enable a local area network VPN connection.

Registering for the My Support function

Complete this task to register for the My Support function.

The My Support function provides proactive notification of code updates through an e-mail address that you specify. My Support automatically notifies you of the latest code updates and how to obtain them. It is highly recommended that you register for My Support.

To access online technical support, go to the following Web site:

<http://www-1.ibm.com/servers/storage/support/disk/ds6800/>

My Support registration provides e-mail notification when new firmware levels have been updated and are available for download and installation. To register for My Support, perform the following steps:

1. Go to the following Web site: <http://www.ibm.com/support/mySupport> and click **register now**.
2. On the My IBM registration Step 1 of 2 page, fill in the required information. Items with an asterisk (*) are required fields. Click **Continue**.
3. On the My IBM registration Step 2 of 2 page, fill in the required information. Items with an asterisk (*) are required fields. Click **Submit**. Click **Continue**.
4. On the Sign in page, enter your IBM ID and password and click **Submit**.
5. On the My support page, click **Edit profile**. Select the information that is required for your profile in the **Products** section.
 - a. From the **Products** list, select **Storage**.
 - b. From the second list, select **Computer Storage**.
 - c. From the third list, select **Disk Storage Systems**.
 - d. From the fourth list, select **System Storage DS6000 series**.
 - e. Select the **System Storage DS6800** check box.
 - f. Click **Add products**.
 - g. Review your profile for accuracy.

6. Click **Subscribe to email**. In the **Documents** list, select **Storage**. Then select the **Please send these documents by weekly email** check box, select the **Downloads and drivers** check box, and select the **Flashes** check box. Click **Update**.
7. In the **Welcome** area, click **Sign out** to end your session.

Activating the Call Home function using the DS CLI

Complete this task to activate the Call Home function using the DS CLI. The Call Home function allows the transmission of operational and error-related data to IBM. It is the capability that allows the storage unit to alert IBM support of machine conditions.

Note: The Call Home function also has a modem call home feature. The DS CLI cannot be used to configure this feature. You must use the DS Storage Manager to configure the modem call home feature.

Before you can activate and test the Call Home function, ensure that the following tasks have been completed:

- You have supplied your required contact information using the **setcontactinfo** command. This consists of the following information:
 - Company name
 - Ship state
 - Ship location
 - Ship country
 - Primary contact phone number
 - Primary e-mail address (required if you are using the call home e-mail notification feature).
- You have supplied the location where e-mail notifications are sent by specifying the SMTP port and IP address through the **setsmtp** command.
- You have set up your account information using the **setplex** command.

Use the **setdialhome** command to enable the Call Home function. You can designate the machine and enable the function on the machine with this command. Use the **testcallhome** command to verify that the Call Home function is active.

Perform the following steps to activate and test the Call Home function for your storage unit. The example commands that are displayed in this task are shown in two formats. The first format provides the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **setdialhome** command to activate the Call Home function. Enter the **setdialhome** command at the dscli command prompt with the following parameters and variables:

```
dscli>setdialhome -action enabled storage_image_ID
```

Example

```
dscli>setdialhome -action enabled IBM.1750-68FA120
```

2. Press Enter. The following message is displayed if the process is successful:

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.1750-68FA120
```

```
The dial home settings were successfully modified
```

3. Issue the **testcallhome** command to verify that the Call Home function is active. Enter the **testcallhome** command at the dscli command prompt with the following parameter:

```
dscli>testcallhome storage_image_ID
```

Example

```
dscli>testcallhome IBM.1750-68FA120
```

4. Press Enter. The following message is displayed if the process is successful:

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.1750-68FA120
```

A test problem record was successfully created.

Activating remote support (real-time only)

Complete this task to initiate a virtual private network (VPN) remote support connection through your local area network or through a dial-up modem.

This task enables you to connect to IBM for support. The modem always initiates the VPN connection if there are phone numbers that are configured on the Configure modem remote support page. If you want to initiate a VPN connection through your local area network, you must unconfigure the modem phone numbers. Follow the steps in “Unconfiguring modem phone numbers (real-time only)” to enable a local area network VPN connection.

Note: Prior to starting the VPN session, notify the support organization that a VPN connection will time out after a period of time with no activity.

Perform the following steps to active remote support:

1. In the navigation, select **Real-time Manager** → **Manage Hardware** → **Storage units**. Select the appropriate storage unit from the table. From the **Select Action** list, select **Activate remote support** and then click **Go**. The Activate remote support page is displayed.
2. Click **Connect**. The connection to IBM support is established. An IBMVPN connection icon is displayed in the system tray of the management console. When you are finished using the VPN connection, right-click on the system tray icon and select **Disconnect** to close the connection.

Note: If the IBMVPN connection icon is not in the system tray of the management console, contact IBM support.

Unconfiguring modem phone numbers (real-time only)

Complete this task to remove modem phone numbers to allow VPN connection through the Internet.

If a modem was previously configured but you now want to create a VPN connection through the Internet, you must first remove all previously completed fields from the Configure Modem Remote Support page. If the modem has been configured with phone numbers to dial, the DS Storage Manager only uses the modem to create a VPN connection. If the modem connection fails for any reason, there can be no IBM remote support connection. The program does not attempt to connect using VPN over the Internet.

Perform these steps to remove modem phone numbers to initiate a virtual private network (VPN) connection through the Internet. IBM can use this connection to perform remote support on your DS6000.

1. In the navigation, under **Real-time manager**, select **Manage hardware** and then select **Storage complexes**. Select **Configure Modem Remote Support** from the **Select Actions** list and click **Go**.
2. Clear all entries from the telephone number fields, leaving the fields empty.
3. Click **OK** to permanently remove the information. The next VPN connection that you initiate will use the Internet.

Chapter 7. Creating an initial logical storage configuration using DS Storage Manager

This section contains high-level logical storage configuration roadmaps that can guide you through the creation of your initial logical storage configuration.

You must perform only one of the following two tasks to create an initial logical storage configuration:

- **Creating an express logical storage configuration:** Follow one of these roadmaps if you want to create a simulated or real-time logical storage configuration where the DS Storage Manager automatically determines the storage allocation details. These tasks guide you through the Express Configuration wizard where you can answer a few basic questions and the DS Storage Manager creates a basic logical storage configuration to match your storage needs.
- **Creating a custom logical storage configuration:** Follow these steps if you want to create a custom simulated or real-time logical storage configuration where you manually determine the storage allocation details. This task can guide you through the pages in the DS Storage Manager that allow you to determine the basic logical storage configuration settings as well as fine-tune the storage allocation details.

Creating an express configuration

This section contains the high-level roadmaps for creating express logical configurations.

You can create the following express logical storage configurations:

- Open systems logical storage configurations
- zSeries logical storage configurations
- iSeries logical storage configurations

Creating open systems volumes using express configuration

Complete this task to quickly configure a storage complex with open systems volumes.

Ensure that you have obtained and applied an Operating Environment License with a maximum storage amount that exceeds the amount of storage that you have configured in your simulated configuration.

To create open systems volumes using express configuration, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Express configuration wizard**. The Express configuration wizard page appears.
2. From the **Select storage unit** list, select the storage unit for the volumes that you are configuring.
3. From the **Select volume type** list, select **Standard Open System (FB)**, and then click **Next**. The Open systems volumes page appears.
4. Select either **RAID 5** or **RAID 10** under **Select RAID type**.

5. Make selections for two of the following three fields and then click **Calculate** configure the amount of storage:
 - Select a value from the **Amount of available storage to use** drop-down list. If you select **User defined GB**, you must enter a value, in GB, in the **User defined GB** field. If you select **User defined %**, you must enter a percentage value in the **User defined %** field.
 - Enter an amount, in GB, in the **Volume size** field to determine the size of the individual volumes. The DS Storage Manager automatically determines the number of volumes that are created.
 - Enter an amount, as an integer, in the **Volume quantity** field to determine the number of individual volumes. The DS Storage Manager automatically determines the size of the volumes that are created.
6. To enable host creation for the open system volumes, select **Create host**.
7. Click **Next**. The Set volume naming page is displayed.
8. To specify a sequence of volume names for the quantity that you created, select **Generate a sequence of nicknames based on the following**. Select the **Use hexadecimal sequence** check box to generate a hexadecimal sequence. See "Using hexadecimal values," on page 253 for conversion information.
9. Select **Verify nicknames are unique** to determine if volume names that are generated on this page are unique.
10. Enter the prefix (alphabetic) that you want for the volumes in the **Prefix** field.
11. Enter the suffix (numeric) that you want for the volumes in the **Numeric suffix** field. Click **Next**. If you selected **Create host**, the Group volumes page appears.
12. Enter a name (up to 16 characters) for the volume group in the **Volume group name** field, select the volumes to include in the volume group under **Select volumes**, and then click **Next**.
 - If you did not select to create a host, the Verification page appears and you can move to the final step.
 - If you selected to create a host, the General host information page appears and you can continue to the next step.
13. Select the host system for the open systems volumes in the **Host system** field, and then enter a nickname (up to 16 characters) for the host in the **Nickname** field.
14. (Optional) Enter a description (up to 256 characters) for the host, and then click **Next**. The Host ports page appears.
15. Enter the number of host ports to configure for the open systems volumes in the **Quantity** field, and then select the host attachment type from the **Type** list.
16. To group the ports that you are configuring and manage them as a single host attachment, select **Group ports to share a common set of volumes**, and then click **Next**. The Define WWPNS page appears.
17. Select or enter the 16-digit WWPNS for each host port that you want to include in the identifier, and then click **Next**. The Assign host to volume group page is displayed.
18. To map the volume group to the host attachment, select **Assign host attachment to volume group**, and then click **Next**.
19. Review the details of the configuration on the Verification page. You can navigate through the Express Configuration pages to make any changes to the

configuration by clicking **Back** or **Next**, or by selecting a specific step in the wizard in the left navigation. When you are satisfied with the details of the configuration, click **Finish**.

Creating zSeries volumes using express configuration

Complete this task to quickly and easily configure a storage complex with zSeries volumes.

Ensure that you have obtained and applied an Operating Environment License with a maximum storage amount that exceeds the amount of storage that you have configured in your simulated configuration.

Note: Before you configure, disable any firewalls, as they might interfere with DS6000 communication.

To create zSeries volumes using express configuration, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Express configuration wizard**. The Express configuration wizard page appears.
2. From the **Select storage unit** list, select the storage unit for the volumes that you are configuring.
3. In the **Select volume type** list, select **zSeries (CKD)**, and then click **Next**. The zSeries volumes page appears.
4. Select either **RAID 5** or **RAID 10** under **Select RAID type**.
5. Select the appropriate model in the **Volume type** list.
6. Select one of the following options to configure an amount of available storage:
 - Select **Amount of available storage to use**. Either select a percentage from the list or enter a value in one of the **User defined** fields representing the space to configure.
 - Select **Volume quantity**, and enter the number of volumes that you want to create.
7. Click **Calculate**. The value for the item that you did not select is automatically displayed.
8. Click **Next**. The LCU/SSID page appears.
9. Select the LCU that you want as the starting LCU in the **Select starting LCU** list.
10. Select the type for the LCU in the **LCU type** list.
11. If you want to change the default SSID, highlight the SSID in the **SSID** field, and type your selection.
12. If the **PAVs** section appears, optionally select **Define number of aliases per base** and enter a number in the **Aliases/base** field.
13. Enter the appropriate number in the **Number of LCUs to create** field, and then select one of the following options:
 - Select **Spread volumes equally across LCUs**.
 - Select **Utilize all addresses in each LCU** to have the quantity of base volumes that are created use all the addresses in the specified number of LCUs.
14. Click **Next**. The Set volume naming page appears.
15. To specify a sequence of volume names for the quantity created, select **Generate a sequence of nicknames based on the following**.

16. Select **Verify nicknames are unique** to determine whether volume names that are generated on this page are unique.
17. Enter the prefix (alphabetic) that you want for the volumes in the **Prefix** field.
18. Enter the suffix (numeric) that you want for the volumes in the **Numeric suffix** field. If you are using hexadecimal values, see "Using hexadecimal values," on page 253 for conversion information. Click **Next**.
19. Review the details of the configuration on the Verification page. You can navigate through the Express Configuration pages to make any changes to the configuration by clicking **Back** or **Next**, or by selecting a specific step in the wizard in the left navigation. When you are satisfied with the details of the configuration, click **Finish**.

Creating iSeries volumes using express configuration

Complete this task to quickly configure a storage complex with iSeries volumes.

Ensure that you have obtained and applied an Operating Environment License with a maximum storage amount that exceeds the amount of storage that you have configured in your simulated configuration.

Note: Before you configure, disable any firewalls, as they might interfere with DS6000 communication.

To create iSeries volumes using express configuration, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Express configuration wizard**. The Express configuration wizard page appears.
2. From the **Select storage unit** list, select the storage unit for the volumes that you are configuring.
3. In the **Select volume type** list, select **iSeries (FB)**, and then click **Next**. The iSeries volumes page appears.
4. Select either **RAID 5** or **RAID 10** under **Select RAID type**.
5. Select either **Protected** or **Unprotected** under **Select volume type**.
6. Select one of the provided iSeries volume sizes from the **Select volume size** field.
7. Select one of the following options to configure an amount of available storage:
 - Select **Amount of available storage to use** and either select a percentage from the list or enter a value in one of the **User defined** fields representing the space to configure.
 - Select **Volume quantity** and enter the number of volumes that you want to create.
8. Click **Calculate**. The value for the item that you did not select is automatically calculated.
9. To enable host creation for the iSeries volumes, select **Create host**.
10. Click **Next**. The Set volume naming page is displayed.
11. To specify a sequence of volume names for the quantity created, select **Generate a sequence of nicknames based on the following**.
12. Select **Verify nicknames are unique** to determine if the volume names that are generated on this page are unique.
13. Enter the prefix (alphabetic) that you want for the volumes in the **Prefix** box.

14. Enter the suffix (numeric) that you want for the volumes in the **Numeric suffix** box. If you are using hexadecimal values, see “Using hexadecimal values,” on page 253 for conversion information. Click **Next**.
15. Enter a name (up to 16 characters) for the volume group in the **Volume group name** box, select the volumes to include in the volume group under **Select volumes** and click **Next**. The General host information page is displayed.
 - If you did not select to create a host, the Verification page is displayed and you can move to the final step.
 - If you selected to create a host, the General host information page is displayed and you can continue to the next step.
16. Select the host system for the iSeries volumes in the **Host system** box, and then enter a nickname (up to 16 characters) for the host in the **Nickname** box.
17. (Optional) Enter a description (up to 256 characters) for the host and click **Next**. The Host ports page is displayed.
18. Enter the number of host ports to configure for the iSeries volumes in the **Quantity** box, and then select the host attachment type from the **Type** list.
19. To group the ports that you are configuring and to manage them as a single host attachment, select **Group ports to share common set of volumes** and click **Next**. The Define WWPNS page is displayed.
20. Select or enter the 16-digit WWPNS for each host port that you want to include in the identifier and click **Next**. The Assign host to volume group page is displayed.
21. To map the volume group to the host attachment, select **Assign host attachment to volume group** and click **Next**.
22. Review the details of the configuration on the Verification page. You can navigate through the Express Configuration pages to make any changes to the configuration by clicking **Back** or **Next**, or by selecting a specific step in the wizard in the left navigation. When you are satisfied with the details of the configuration, click **Finish**.

Creating a custom logical storage configuration

Complete this task to create a simulated or real-time custom logical configuration where you manually determine the storage allocation details. You must complete all of the steps and substeps in this task to create your configuration.

You must have installed and configured the DS Storage Manager before you can create your custom logical configuration. You also must have obtained and applied an Operating Environment License (OEL) with a maximum storage amount that exceeds the amount of storage that you intend to configure. In addition, your network must be configured so that it can support all the components and functions that you will use with your storage unit.

If you are working in simulated mode, you must create a new storage configuration or open an existing storage configuration before completing this task. Follow the steps in “Creating a configuration file (simulated only)” on page 91 to manage your configuration files.

You can create either a simulated or real-time custom logical configuration by following the steps in this task. A real-time configuration allows you to manage physical and logical configurations from existing storage complexes, and storage servers over your network in real time. A simulated configuration allows you to create or import a new *simulated* instance of your DS6000 series.

Note: Before you configure, disable any firewalls, as they might interfere with DS6000 communication

Return to the next step of this task after completing each of the subtasks below.

1. Create your arrays. To create an array, you must complete the subtask in “Creating arrays.”
2. Create your ranks. To create a rank, you must complete the subtask in “Creating ranks” on page 169.
3. Create your extent pools. To create an extent pool, you must complete the subtask in “Creating extent pools” on page 170.
4. Perform one of the following tasks for either your zSeries or open systems host:
 - If you are working with a zSeries host, you must create zSeries LCUs and zSeries volumes.
 - a. Create zSeries LCUs. To create zSeries LCUs, you must complete the subtask in “Creating zSeries LCUs” on page 171.
 - b. Create zSeries volumes. To create zSeries volumes, you must complete the subtask in “Creating zSeries volumes” on page 172.
 - If you are working with an open systems host, you must create open systems volumes and open systems volume groups.
 - a. Create open systems volumes. To create open systems volumes, you must complete the subtask in “Creating open systems volumes” on page 173.
 - b. Create open systems volume groups. To create open systems volume groups, you must complete the subtask in “Creating open systems volume groups” on page 174.
5. Reconfigure your host system to start sending I/O to the DS6000. To reconfigure your host system, you must complete the subtask in “Enabling host system I/O” on page 176.

Note: If you are working in simulated mode, you must apply the configuration before you can complete this step. To apply the configuration, complete the subtask in “Applying a configuration (simulated only)” on page 175.

Creating arrays

Complete this task to create arrays, either automatic or custom, and to specify their RAID and rank attributes.

Perform the following steps to create arrays:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Arrays**. On the Arrays main page, in the **Select storage unit** drop-down list, select a storage unit.
2. From the **Select Action** list, select **Create...**, and then click **Go**. The Create Array – Definition method page is displayed.
3. In Create Array – Definition method, choose to either create your arrays automatically, where the DS Storage Manager specifies the array sites, or to create custom arrays where you specify the array sites. Click **Next**.
 - If you chose to create your arrays automatically, in Create Array – Array configuration (Auto), specify the quantity and RAID type of the arrays that you are creating. You can optionally select **Create an 8 disk array** to create an array that contains two, four-disk array sites with the same specifications. Click **Next** after you have made your selections. The Create Array – Add arrays to ranks page is displayed.

- If you chose to create custom arrays, perform the following steps:
 - a. In Create Array – Array configuration (Custom), specify the RAID type for the arrays that you are creating and select the array site numbers that you want to assign to the arrays. You can optionally select **Create an 8 disk array** to create arrays that contain two, four-disk array sites with the same specifications. Click **Next** after you have made your selections.
 - b. If you selected to create an 8 disk array, the Create Array: Real-time – Second array-site selection page is displayed. Choose the second array site for each array that was selected in the Create Array – Array configuration (Custom) page. Secondary array sites cannot be shared among arrays. Click **Next**. The Create Array – Add arrays to ranks page is displayed.

Note: When configuring two 8 disk arrays (one Raid 5, one Raid 10), you must create the Raid 10 array first and make sure that it is being created on the array sites that contain spares. Then create the Raid 5.

4. In Create Array – Add arrays to ranks, you can optionally choose to add the arrays that you are creating to ranks. If you choose to add your arrays to ranks, you must specify the storage type, FB or CKD, for the ranks that you are creating. Click **Next**. The Create Array – Verification page is displayed.
5. In Create Array – Verification, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and then specify the correct values. Otherwise, click **Finish** to complete the array creation process.

If you are creating a new logical storage configuration, return to “Creating a custom logical storage configuration” on page 167 and complete the rest of the steps.

Creating ranks

Complete this task to create ranks and to specify their arrays and extent pools.

Note: If you are building a configuration in simulated mode or you are importing a real-time configuration into the simulated configurator, extent allocation per rank in an extent pool might not be the same. This means that the rank extent allocations that are created in simulated mode or imported as part of a real-time configuration might not match the allocation when you compare applied and imported configurations to real-time configurations at the rank level, although the extent usage is equivalent. The real-time configuration extent allocation at the rank level will be configured appropriately.

Perform the following steps to create ranks:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Ranks**. On the Ranks main page, in the **Select storage unit** drop-down list, select a storage unit.
2. From the **Select Action** list, select **Create...**, and then click **Go**. The Create Rank – Select array for rank page is displayed.
3. In the **Select** column, select an array. Then click **Next**. The Create Rank – Define rank properties page is displayed. The **Rank number** is provided by default.

Note: If there are no available arrays for selection, you can create one by clicking the **Create new array** button. When you have finished creating an array, you are returned to this page, and the array that you created is available for selection.

4. Specify the **Storage type** and then click **Next**. The Create Rank – Select extent pool page is displayed.
5. Optionally, you can select one extent pool. In the **Select** column, select an extent pool. Then click **Next**. The Create Rank – Verification page is displayed.

Note: If there are no available extent pools for selection, you can create one by clicking the **Create new extent pool** button. When you have finished creating an extent pool, you are returned to this page, and the extent pool that you created is available for selection.

6. In Create Rank – Verification, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and then specify the correct values. Otherwise, click **Finish** to complete the rank creation process.

If you are creating a new logical storage configuration, return to “Creating a custom logical storage configuration” on page 167 and complete the rest of the steps.

Creating extent pools

Complete this task to create extent pools, automatic or custom, and to specify extent pool parameters.

Each extent pool is either CKD or FB. An extent pool is associated with either server 0 or server 1. You need a minimum of 2 extent pools, one for each server, to make full use of the resources. An extent pool is divided into equal extents of the following size:

- FB - 1GB
- CKD - 1113 cylinders

Your maximum number of extent pools cannot exceed the number of ranks. It is recommended that you put only ranks with the same capacity and speed into the same extent pool. RAID types cannot be intermixed in an extent pool. For performance management, you can create extent pools of a single rank.

Perform the following steps to create extent pools:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Extent pools**. On the Extent pools main page, in the **Select storage unit** drop-down list, select a storage unit.
2. From the **Select Action** list, select **Create...**, and then click **Go**. The Create Extent Pool – Define properties page is displayed.
3. Choose either to create your extent pool automatically based on storage requirements (where the DS Storage Manager automatically selects your arrays and ranks) or to create a custom extent pool, where you decide which arrays and ranks are assigned to the extent pools.
 - If you chose to create your extent pools automatically, in Create Extent Pool – Define requirements, provide a nickname for your extent pool, select the storage type (either FB or CKD) and RAID type, and enter the required amount of storage. You can optionally select to use any existing unassigned

arrays and ranks if you want the DS Storage Manager to use arrays and ranks that already exist. Click **Next**. The Create Extent Pool – Reserve storage page is displayed.

Note: An extent pool uses all of the available storage in arrays or ranks that make up the extent pool, even if you specify a storage amount that is less than the entire array or rank storage amount. For example, if you have a 385 GB rank and need 100 GB of storage, your extent pool will contain 385 GB of storage. However, if you have two ranks with 100 GB of storage in each rank and you need 120 GB in your extent pool, your extent pool will contain two ranks with 200 GB of storage.

- If you chose to create a custom extent pool, perform the following steps:
 - a. In Create Extent Pool – Define properties, provide a nickname for your extent pool, then select the storage type (either FB or CKD), RAID type, and server that are assigned to the extent pool. Click **Next** after you have made your selections.
 - b. In Create Extent Pool – Select ranks, select the rank that you want to assign to the extent pool or create a new rank to assign to the extent pool. Click **Next**. The Create Extent Pool – Reserve storage page is displayed.
- 4. In the Create Extent Pool – Reserve storage page, optionally enter the amount of storage that you want the DS Storage Manager to reserve for the extent pool. Click **Next**. The Create Extent Pool – Verification page is displayed.
- 5. In the Create Extent Pool – Verification page, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back**, as appropriate, to return and to specify the correct values. Otherwise, click **Finish** to complete the extent pool creation process.

If you are creating a new logical storage configuration, return to “Creating a custom logical storage configuration” on page 167 and complete the rest of the steps.

Creating zSeries LCUs

Complete this task to create zSeries logical control units (LCUs) and to specify their attributes and properties.

To create zSeries LCUs, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **zSeries** → **LCUs**. On the LCUs main page, from the **Select storage unit** list, select a storage unit.
2. From the **Select Action** list, select **Create...**, and then click **Go**. The Create LCU – Select from available LCUs page is displayed.
3. Select one or more LCUs from the list of those available. Click **Next** to continue. The Define LCU properties page is displayed.

Note: If you are creating LCUs to associate with a specific extent pool, you must select LCU ID numbers that match, as even or odd, the server number that was specified for the extent pool when it was created. For example, an extent pool that is associated with Server 1 can be associated only with LCUs that have odd numbers.

4. In the Define LCU properties page, define the parameters for the selected LCUs.
 - a. You can change the default SSID by highlighting and typing over it.
 - b. Specify the LCU types and establish the timeout times in seconds.

5. Click **Next** to continue. The Verification page is displayed.
6. Use the Verification page to review the established attributes and verify that they are correct. If the attributes and values are not correct, click **Back**, as appropriate, to return and to specify the correct values. Otherwise, click **Finish** to complete the LCU creation process.

If you are creating a new logical storage configuration, return to “Creating a custom logical storage configuration” on page 167 and complete the rest of the steps.

Creating zSeries volumes

Complete this task to create zSeries volumes and to specify their attributes and properties.

Only certain LCUs are available for an extent pool. To be available, the LCU numbers must match in parity—even or odd—with the server number that is associated with the extent pool. For example, given an extent pool that was associated with Server 1, which is odd, you will have LCUs for selection in the Define volume characteristics page only if there are LCUs that are also odd. If none appear, you must create LCUs by assigning odd numbers to them in the Create zSeries LCUs wizard.

To create zSeries volumes, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **zSeries** → **Volumes – zSeries**. On the Volumes – zSeries main page, from the **Select storage unit** list, select a storage unit. In the **Select LCU** drop-down list, select a zSeries LCU.
2. From the **Select Action** list, select **Create...**, and then click **Go**. The Create zSeries volume – Select extent pool is displayed.
3. Select the extent pools that you want to use to create your volumes or click **Create new extent pool** to create new CKD extent pools that you can then use to create your zSeries volumes. Click **Next**. The Create zSeries volume – Define base volume characteristics page is displayed.
4. In Create zSeries volume – Define base volume characteristics, specify the volume type and LCU number that are associated with the zSeries volumes. Click **Next**. The Create zSeries volume – Define base volume properties page is displayed.

Note: LCUs are displayed for selection only if their numbers match in parity, even or odd, with the server number that was selected for the extent pool. If no LCUs are displayed, you must create them with the correct parity before you can continue.

5. In Create zSeries volume – Define base volume properties, select an addressing policy and if necessary, supply a maximum LCU address. Enter the quantity of base volumes. If you selected a custom volume type on the Define base volume characteristics page, you have the option to specify a size in cylinders. You also must specify the quantity of LCUs, the base start address, and a sorting order of Ascending or Descending. The Available storage table shows the amount of storage that you can work with while defining your volume properties. Click **Next**. The Create zSeries volume – Create volume nicknames page is displayed.
6. In Create zSeries volume – Create volume nicknames, choose if you want to generate a sequence of nicknames based on the specified prefix and suffix and if you want to use hexadecimal values. If you are using hexadecimal values, see

- “Using hexadecimal values,” on page 253 for conversion information. Click **Next**. The Create zSeries volume – Define alias assignments page is displayed.
7. In Create zSeries volume – Define alias assignments, select the LCUs for which you want to define the aliases. Then specify the starting address and address sorting order. Provide the values for the number of aliases per number of volumes. For example, if you enter 1 in the **Aliases** field, and 4 in the **Per volumes field**, you are assigning one alias name for every four volumes. Click **Next** to continue. The Create zSeries volume – Verification page is displayed.
 8. Use the Verification page to review the attributes and verify that they are correct. If the attributes and values are not correct, click **Back**, as appropriate, to return and specify the correct values. Otherwise, click **Finish** to complete the volume creation process.

If you are creating a new logical storage configuration, return to “Creating a custom logical storage configuration” on page 167 and complete the rest of the steps.

Creating open systems volumes

Complete this task to create open systems volumes and to specify their attributes and properties.

Note: Before you configure, disable any firewalls, as they might interfere with DS6000 communication.

To create open systems volume groups, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Open systems** → **Volumes – Open systems**. On the Volumes – Open systems main page, in the **Select storage unit** list, select a storage unit.
2. Select a volume. From the **Select Action** list, select **Create...**, and then click **Go**. The Create Volume – Select extent pool page is displayed.
3. You must select an extent pool for the target volumes. After you select the extent pool, click **Next**. The Create Volume – Define volume characteristics page is displayed.

Note: If there are no available extent pools available for selection, you can create one by clicking the **Create new extent pool** button. When you have finished creating an extent pool, you are returned to this page, and the extent pool that you created is available for selection.

4. In Create Volume – Define volume characteristics, define the volume type for the new volumes. You can select any number of volume groups from the **Select volume groups** list to associate with the new volumes. After you define the volume characteristics, click **Next**. The Create Volume – Define volume properties page is displayed.
5. Define the volume properties. To calculate the maximum size, enter a value in the **Quantity** field and click the **Calculate max size** button. To calculate the maximum quantity, enter a value in the **Size** field and click the **Calculate max quantity** button. The **Extent size** and **Available Storage Capacity in Extent Pool** fields are displayed in the units that correspond to the selected volume type.

If you click the **Calculate max size** button, the **Size** field is populated with the calculated value. You can overwrite the value and enter an integer value for a lesser size. If you selected one of the iSeries volume types on the previous

Define volume characteristics page and Decimal GB (10^9 bytes) for the capacity units, the values here include 8.56, 17.54, 35.16, 36.00, 70.56, 141.12, and 282.25. If you selected one of the iSeries volume types on the previous Define volume characteristics page and Binary GB (2^{30} bytes) for the capacity units, the values here include 8.00, 16.34, 32.75, 33.53, 65.72, 131.44, and 243.80.

If you select the **Select LSSs for volumes** checkbox, you are required to select from the list of available LSSs for these open systems volumes.

After you define the volume properties, click **Next**. The Create Volume – Create volume nicknames page is displayed.

6. In Create Volume – Create volume nicknames, you can optionally create one or more nicknames that are based on entries in the **Prefix** or **Suffix** fields. If you do not specify nicknames, only a volume number is created. Click **Next**. The Create Volume – Verification page is displayed.

Tip: If you plan to create volume groups, you can use a unique and meaningful nickname that can help you easily find the volumes that you want to include in a volume group. For example, if you are creating multiple volumes of data for a specific department, you can make the nickname prefix an abbreviation of that department's name. Then you can use a predetermined range of numbers in the nickname suffix to identify each individual volume.

7. In Create Volume – Verification, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and to specify the correct values. Otherwise, click **Finish** to complete the volume creation process.

If you are creating a new logical storage configuration, return to “Creating a custom logical storage configuration” on page 167 and complete the rest of the steps.

Creating open systems volume groups

Complete this task to create open systems volume groups and to specify their properties and parameters.

Volume groups provide a volume masking function. When you create a volume, you can assign it to one or more volume groups. Open systems server attachments are assigned to volume groups to control access to the volumes. There can be up to 8320 volume groups. Any host port (or host port group) can access only one volume group. A volume group can be accessed by multiple servers if they have the same block size and address. If there is a requirement for volume sharing, for example in a clustering environment, then a volume can be a member of multiple volume groups, or multiple servers can be in the same volume group. In both cases, data integrity is the responsibility of host software.

To create open systems volume groups, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Open systems** → **Volume groups**. On the Volume groups main page, in the **Select storage unit** drop-down list, select a storage unit.
2. From the **Select Action** list, select **Create...**, and then click **Go**. The Create Volume Group – Define volume group properties page is displayed.
3. Define the properties. In the **Nickname** field, you can either keep or modify the default nickname. In **Accessed by host types**, select the appropriate host type. All other valid compatible host types will be automatically selected. Both the

Nickname and **Accessed by host types** fields are required. After you define the properties, click **Next**. The Create Volume Group – Select host attachments page is displayed.

4. In Create Volume Group – Select host attachments page, you can optionally select the host attachment for the specified volumes. After selecting the host attachment, click **Next**. The Create Volume Group – Select volumes for groups page is displayed.

Note: If there are no host attachments available for selection, you can create one by clicking the **Create new host attachment** button. When you have finished creating a host attachment, you will be returned to this page, and the host attachment that you created will be available for selection.

5. Select the volumes for the volume group. At least one volume must be selected in the table. The volumes listed should be compatible with the host types selected in the previous two pages. After selecting the volumes for the volume group, click **Next**. The Create Volume Group – Verification page is displayed.

Note: If there are no volumes available for selection, you can create one by clicking the **Create new open systems volume** button. When you have finished creating a volume, you will be returned to this page, and the volume that you created will be available for selection.

Tip: If you created volumes with the same nickname prefix, you can change the view to show only those volumes that you want to select either by sorting the nickname column or by creating a filter.

6. Use the Create Volume Group – Verification page to review the attributes and verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and to specify the correct values. Otherwise, click **Finish** to complete the volume group creation process.

Applying a configuration (simulated only)

Complete this task to select, authenticate, and apply a storage unit configuration.

Ensure that you have obtained and applied an Operating Environment License with a maximum storage amount that exceeds the amount of storage that you have configured in your simulated configuration.

The source storage unit must have defined a logical configuration and hosts before you can apply a configuration.

To apply a configuration in simulated mode, perform the following steps:

1. In the navigation, select **Simulated manager** → **Manage hardware** → **Storage units**.
2. From the **Select Action** list, select **Apply configuration...**, and then click **Go**. The Apply Configuration — Select application method page is displayed.
3. In the Select application method page, specify the method with which to apply the configuration.
 - a. If you choose **Select from a list of storage complex**, the Select storage complex page of this wizard is displayed with the storage complexes in your simulated environment. Click **Next** to continue and go to step 3.
 - b. If you choose **Import new storage complex** the Import storage complex wizard is displayed. Once you finish with the wizard, the Select storage unit page is displayed with the storage units from the imported storage

complex. You must be connected to the network to import the storage complex. Click **Next** to continue and go to step 5.

- c. If you choose the **Apply configuration without importing storage complex** option, the Authenticate page is displayed. Click **Next** to continue and go to step 4.
4. Use the Select storage complex page to connect directly to a storage unit. Click **Next** to continue. The Authenticate page is displayed.
5. In the Authenticate page, connect to and authenticate a storage complex by defining management console properties. You must provide a user ID and password to complete the authentication. Click **Next** to continue. The Select storage unit page is displayed.
6. On the Select storage unit page, you connect directly to a storage unit. Specify the values as appropriate and click **Next**.
7. Use the Verification page to review the attributes and verify that they are correct.
8. If the attributes and values are not correct, click **Back**, as appropriate, to return and specify the correct values. Otherwise, click **Finish** to complete the apply configuration process.

If you are creating a new logical storage configuration, return to “Creating a custom logical storage configuration” on page 167 and complete the rest of the steps.

Enabling host system I/O

Complete this task to perform the high-level steps that are required to enable your host to send I/O to the DS6000.

You must complete your logical storage configuration before you attempt to send host I/O to the DS6000.

This task contains the high-level tasks that you must complete before you can successfully send host I/O to your DS6000. Instructions and details for each individual step can be found in the host system documentation or the *IBM System Storage DS6000 Host Systems Attachment Guide*.

1. Verify and update the supported fabric switch firmware.
2. Create your storage area network (SAN) zones and assign them to the appropriate zone sets.
3. Verify and update the host adapter firmware on the host system.
4. Install, verify, and update the host failover driver (SDD) on the host system.
5. If necessary, reconfigure or reboot the host system.
6. Verify that the host system can access the DS6000 logical volumes that are assigned to it.

If you are creating a new logical storage configuration, return to “Creating a custom logical storage configuration” on page 167 and complete the rest of the steps.

Chapter 8. Creating a logical storage configuration using the DS command-line interface

This section contains the command-line interface (CLI) roadmaps and tasks that you must complete to create a logical storage configuration.

You can create the following two types of logical storage configurations:

- **Fixed block logical storage configurations:** Fixed block logical storage configurations are required for open systems host attachment.
- **Count key data (CKD) logical storage configurations:** CKD logical storage configurations are required for zSeries host attachment.

Configuring new fixed block storage using the DS CLI

This section describes how you can configure new fixed block storage within a storage unit by using the command-line interface.

Before you begin, you must be logged into the DS CLI application in interactive command mode. You must also be connected to a storage unit that is used for open systems host system storage.

The creation of the fixed block storage configuration is described first. The configuration of the storage unit SCSI host ports to enable access to fixed block storage is described second. You can run these two basic steps in the reverse order, but it is better to create storage configurations first, thereby creating the media to back up configuration data that is not related to the storage configuration.

Configuring new fixed block storage involves the following processes:

- Creating fixed block extent pools
- Creating arrays
- Creating and associating ranks with extent pools
- Creating fixed block volumes
- Creating fixed block volume groups
- Configuring fibre-channel I/O ports
- Creating SCSI host port connections

Note: All the examples provided in the described tasks are based on the premise of using the interactive mode of DS CLI. If you were processing many transactions, you would likely use the script mode to process your transactions.

Creating extent pools for fixed block volumes using the DS CLI

Complete this task to create fixed block volume extent pools. This is the first step in configuring new fixed block storage. You can use the DS CLI commands to create extent pools for fixed block volumes.

Creating the extent pools before the arrays and ranks saves a processing step. When you create the new ranks, you can assign them to existing extent pools.

Otherwise, you must modify each rank object to complete the extent pool ID assignment after the extent pools have been defined.

Each extent pool is defined with the rank group of 0 or 1 and a storage type of **fb**. You must define one extent pool for each rank group and storage type combination. This means that you must make a minimum of two extent pools for a storage unit that contains fixed block storage: one fixed block extent pool per rank group.

Extent pools that are defined for rank group 0 or 1 are assigned an even- or odd-numbered extent pool ID, respectively. Even-numbered extent pools are managed by storage server ID 0. Odd-numbered extent pools are managed by storage server ID 1. Each rank is assigned to one extent pool; therefore, storage server workload is affected by the rank assignments to even- and odd-numbered extent pool IDs. It is better to evenly distribute rank and extent pool allocations in order to keep the storage server workloads balanced.

You can create more than the minimum number of extent pools. For example, you can define unique extent pools for each RAID type (5 or 10) that is configured in a storage image. Or, you can define and name extent pools according to the host system attachments that access the volumes that are created from extent pool extents. You can have the same number of extent pools as ranks.

i5/OS considerations

i5/OS supports only specific volume sizes and these might not be an exact number of extents. i5/OS volumes are defined in decimal gigabytes (10⁹ bytes). You can use the following table when you are creating the logical volumes for use with i5/OS. You will notice that in almost every case, the i5/OS device size does not match a whole number of extents, so some space can be wasted for you specific configuration.

Protected Model Type	Unprotected Model Type	i5/OS Device size (decimal gigabytes)	Number of LBAs	Extents	Unusable space (binary gigabytes)	Usable space%
xxxx-A01	xxxx-A81	8.5	16 777 216	8	0.00	100.00
xxxx-A02	xxxx-A82	17.5	34 275 328	17	0.66	96.14
xxxx-A05	xxxx-A85	35.1	68 681 728	33	0.25	99.24
xxxx-A04	xxxx-A84	70.5	137 822 208	66	0.28	99.57
xxxx-A06	xxxx-A86	141.1	275 644 416	132	0.56	99.57
xxxx-A07	xxxx-A87	282.2	551 288 832	263	0.13	99.95

Note: Only Ax2, Ax4 and Ax5 models are supported as external LSU LUNs.

Use the **lsextpool** and **mkextpool** commands to create the fixed block extent pools. You must be logged into the DS CLI application and connected to the storage unit that will be used for open systems host system storage.

Perform the following steps to create the fixed block extent pools. The example commands displayed in this task are shown in two formats. The first format shows

the type of information the command requires. The second format provides the command with declared values for the variables

1. Issue the **mkextpool** command to create the fixed block extent pool for rank group 0. Enter the **mkextpool** command at the dscli command prompt with the following parameters and variables:

```
dscli>mkextpool -dev storage_image_ID -rankgroup [0 | 1]
-stgtype fb extent_pool_name
```

Example

```
dscli>mkextpool -dev IBM.1750-68FA120 -rankgrp 0 -stgtype fb P0
```

where *P0* represents the extent pool name that you assign. This name can be 16 double-byte characters.

2. Press Enter. A successful process displays the following message:

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.1750-68FA120
```

```
Extent pool P0 successfully created.
```

Note: The unique name that you assigned to the extent pool does not display in the process message. However, when you issue the **lsxtpool** command, the extent pool name is displayed.

3. Repeat Step 1 for each extent pool that you want to create. Try to evenly distribute rank and extent pool allocations in order to keep the storage server workloads balanced.
4. Verify the extent pool assignments by issuing the **lsxtpool** command when you are done creating the extent pools. Use the **-l** parameter to display a full report for the extent pools that are assigned to the storage unit. Enter the **lsxtpool** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsxtpool -dev storage_image_ID -l
```

Example

```
dscli>lsxtpool -dev IBM.1750-68FA120 -l
```

Creating arrays for fixed block volumes using the DS CLI

Complete this task to create arrays using the DS CLI commands.

The machine type 1750 contains at least one storage enclosure, with a minimum of four DDMs.

The DDMs of a storage enclosure are partitioned into array sites. A machine type 1750 array site consists of four DDMs in one storage enclosure of a storage enclosure pair, with two-to-eight (four DDM) array sites per storage enclosure pair. All storage enclosure pairs must have identical capacity, rpm, and interface characteristics, and an interface to a common DA pair.

The creation of arrays is based on the array sites that are associated with the storage unit. Use the **lsarraysite** and **mkarray** commands to create the arrays.

You want to make an array from 1 or 2 array sites. An array inherits the characteristics of its parent array sites and is given a RAID type attribute (5 or 10). A 1750 array object of RAID type 5 or 10 is made from one or two (4 DDMs) array sites. The status of the array is “unassigned” until the array is assigned to a rank.

Perform the following steps to create an array from unassigned array sites:

1. Issue the **lsarraysite** command to view a list of array site IDs for all installed array sites. Review those arrays that are designated with the state of unassigned. Enter the **lsarraysite** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsarraysite -dev storage_image_ID -state unassigned
```

Note: If this is your first time creating fixed block volumes, all the arrays are displayed with a state of unassigned.

2. Press Enter. A report of unassigned array sites is displayed. Use the list to identify unassigned array site capacity, rpm, and device adapter (DA) pair attributes. Record the RAID type for each array site.

3. Issue the **mkarray** command to create an array from either one or two array sites with the status "unassigned". Enter the **mkarray** command at the dscli command prompt with the following parameters and variables:

```
dscli>mkarray -dev storage_image_ID -raidtype [5 | 10] -arsite array_site
```

Consider the following when you create the arrays:

- Specify one or two array sites with identical capacity, rpm, interface, and DA pair attributes.
 - The new array inherits the capacity, rpm, interface, and DA pair characteristics of its parent array sites.
 - The state of the array remains unassigned until it is assigned to a rank.
4. Repeat Step 3 until all unassigned array sites have been assigned to an array.
 5. Verify that the array-to-array site assignment is recognized and complete by issuing either the **lsarray** or **lsarraysite** command with the **-l** parameter.

Creating a rank using the DS CLI

Complete this task to create a rank using the DS CLI commands. A rank is a logically contiguous storage space that is made up of one array. You can assign a rank to every unassigned array.

A rank inherits the characteristics, including the RAID type, of its parent array and is given a storage type attribute of either FB (fixed block) or CKD (count key data). The rank configuration state is unassigned until it is assigned to an extent pool. An "unassigned" rank is not associated with either rank group 0 or 1. Any unassigned rank can be assigned to an extent pool that is associated with either rank group 0 or 1.

Note: You can assign a rank to an unassigned array and also assign the rank to an extent pool at the same time if you have already created the extent pools and the arrays. Creating extent pools first saves a step in the configuration.

Use the **lsarray**, **mkrank**, and **lsrank** commands to assign a rank to each unassigned array. You must be logged into the DS CLI application and connected to the storage unit that will be used for open systems host system storage.

To create ranks, perform the following steps:

1. Ensure you have a list of the unassigned arrays for which ranks must be assigned. Issue the **lsarray** command to obtain this list if you do not already have it. Enter the **lsarray** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsarray -dev IBM.1750-68FA120 -state unassigned
```

2. Issue the **mkrank** command to assign a rank to rank group 0 or 1 according to the rank group number of the assigned extent pool ID. Enter the **mkrank** command at the dscli command prompt with the following parameters and variables:

```
mkrank -dev IBM.1750-68FA120 -array A44 -stgtype fb
-extpool P1
```

Notes:

- a. You can specify either the **-wait** or the **-extpool** parameter when you use the **mkrank** command. Either of these parameters allows you to be notified if the rank configuration has failed for any reason.
 - b. If you use the **-wait** parameter, you cannot issue other commands until the entire transaction has processed.
3. Press Enter to display a report of rank assignments for your entire storage unit. Because the process of creating the rank involves formatting drives, it could take some time before the process finishes. If you want to check on the process, you can issue the **lsrank** command from a different DS CLI session. A successful process displays the following type of message:
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Device: IBM.1750-68FA120

Rank IBM.1750-68FA120/R44 successfully created.
 4. Repeat Step 2 until all unassigned arrays are assigned a rank and an extent pool.
 5. Issue the **lsrank** command to verify that ranks and extent pools have been assigned. Enter the **lsrank** command at the dscli command prompt with the following parameters and variables:
dscli>lsrank -dev IBM.1750-68FA120 -1
 6. Press Enter to display a report of the rank assignments for your entire storage unit.

Correcting a rank-related configuration error

Complete this task to correct a rank-related configuration error.

There might be occasions when you are using the **mkrank** command to create ranks, but the transaction fails. You might not be aware of the failure until you run the **lsrank** or **showrank** command to check the status of the ranks that you have created.

The **lsrank** or the **showrank** commands display reports that includes a *state* category. The state category reports on the current state of the rank. One of the state codes is **configuration error**.

A state of **configuration error** specifies that a rank configuration process has not completed successfully. This state reflects an internal error condition and is not an indication that there was a user input error.

You might want to gather additional information about what caused the error, which can help you determine how to correct it. Generally, to correct this error state, you must delete the designated rank configuration and submit a new transaction request.

Perform the following steps to obtain additional information about the configuration error and to correct this error condition.

1. Obtain additional information about the transaction by implementing one of the following methods:

- Add the **-v** (verbose) command flag to your **mkrank** command and reissue the command for the transactions that show the configuration error designation.

Note: You can also turn on the verbose mode in your profile file and reissue the command.

Designating the verbose mode allows the display of extra output that includes the error code that is generated when the create rank transaction fails.

- Add the **-extpool** parameter to your **mkrank** command and reissue the command for the transactions that show the configuration error.

You might consider using this parameter if you have not yet assigned your ranks to the extent pools. If the transaction fails, a message states the reason for a failure.

2. Issue the **rmrank** command to delete the designated rank configurations if you do not want to obtain additional information about what caused the configuration error.

Note: In the majority of instances, this is the only method for correcting a configuration error.

Creating fixed block volumes using the DS CLI

Complete this task to create fixed block volumes.

You must have completed the following tasks before you can create your fixed block volumes:

- Created your extent pools.
- Created your arrays.
- Created and assigned your ranks.

Perform the following steps to create fixed block volumes:

1. View your list of fixed block extent pool IDs and determine which extent pool IDs that you want to use as the source for the fixed block logical volumes. You obtained this list when you first created your extent pools. If this list is not available, issue the **lsextpool** command to obtain the list of extent pool IDs. Enter the **lsextpool** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsextpool -dev IBM.1750-13AAD7A -stgtype fb -l
```

Extent pool attributes determine the size and quantity of volumes that can be created. The extent pool ID (even/odd) indicates the storage server (0|1), which dictates that the LSS ID component of the volume ID must be an even or an odd number.

2. Issue the **lsaddressgrp** command to find unassigned and available address groups. Enter the **lsaddressgrp** command at the dscli command prompt with the following parameters and variables:

```
dscli> lsaddressgrp -dev IBM.1750-68FA120 -l
```

An address group refers to a group of LSSs. Up to 16 LSSs can be grouped into one address group. All LSSs in an address group must be of the same format (CKD or fixed block).

Note: If this is your first time creating fixed block volumes, all the address groups are displayed with a state of "unassigned".

3. Analyze the address group list to determine which LSSs can be used to make fixed block volumes.

Consider the following conditions when doing your analysis:

- If the address group list is empty, then all address groups are available to be defined (0 - 3).
- If an undefined address group is used to create new fixed block volumes, select the lowest numbered address group.
- If you are adding new fixed block volumes to an existing fixed block address group, use the **lslss** command to identify LSSs that are already defined in the target address group.

4. Issue the **mkfbvol** command to create fixed block volumes for the specified LSS. Enter the **mkfbvol** command at the dscli command prompt with the following parameters and variables:

```
dscli>mkfbvol -dev IBM.1750-68FA120 -extpool P1  
-name finance#d -cap 8.6 0100-010f
```

Consider the following conditions with regard to the command example in this step:

- All volumes will have the same type and capacity attributes.
- The **-extpool** parameter identifies a fixed block extent pool containing available data extents.
- The **-name** parameter allows you to assign an easy-to-use label or nickname to the volume. The volume name parameter can include a wild card (#d or #h) that inserts a decimal or hexadecimal volume ID value into the volume name.

Note: The decimal designation does not apply to the volume ID number or the number of volumes that were created by the command. It only applies to the unique name that you have assigned. Also, when you process this command, the volume name that you have assigned does not appear in the confirmation message. To view the volume name that you have assigned, issue the **lsfbvol** or **showfbvol** command.

- The **-cap** (capacity) parameter is 8.6 GB. The default is binary GB where 1 GB = 1 073 741 824 (2³⁰ bytes)
- The example provides a range of numbers (0100 - 010F) for the number of volumes to be created. Because volumes are created using the hexadecimal numbering system, the range in the example creates 16 volumes. The actual number of volumes that can be created is 255 per LSS based on the following criteria:
 - The volume ID is a 32 bit number that can be represented as 4 hexadecimal digits in the form of XYZZ where X is the address group (0 - 1), XY together is the logical subsystem number 00 - 1E, and ZZ is the volume number (00 - FF).
 - DS6000 has a 16 384 volume address space that is partitioned into 64 logical subsystem (LSS) units, where each LSS contains 256 logical volume numbers. The 64 LSS units are assigned to one of 4 address groups, where each address group contains 16 LSSs, or 4096 volume addresses. All of the LSSs in one address group must be of the same type (CKD or fixed block).

5. Repeat step 4 for all of the required logical volumes for each LSS.

6. Issue the **lsfbvol** command to display a report you can use to confirm the status of your fixed block volumes. Enter the **lsfbvol** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsfbvol -dev IBM.1750-68FA120 -l -volgrp V2,V20
```

Note: It is possible that the report will display that there was a configuration error that is associated with one or more of your **mkfbvol** transactions.

Correcting a fixed block configuration error

Complete this task to correct a fixed block volume configuration error.

There might be occasions when you are using the **mkfbvol** command to create fixed block volumes, but the transaction fails. You might not be aware of the failure until you run the **lsfbvol** or the **showfbvol** command to check the status of the volumes that you have created.

The **lsfbvol** or the **showfbvol** commands display reports that includes a *configstate* category. The configuration state category reports on the current state of the rank. One of the configuration state codes is **configuration error**.

A status of **configuration error** specifies that the configuration process did not complete successfully. This state reflects an internal error condition and is not an indication that there was a user input error.

You might want to gather additional information about what caused the error, which can help you determine how to correct it. Generally, to correct this error state, you must delete the designated volume configuration and submit a new transaction request.

Perform the following steps to obtain additional information about the configuration error and to correct this error condition.

1. Add the **-v** (verbose) command flag to your **mkfbvol** command, and reissue the command for the transactions that show the configuration error designation.

Note: You can also turn on the verbose mode in your profile file, and reissue the command.

If you designate the verbose mode, the display of extra output includes the error code that is generated when the create rank transaction fails.

2. Issue the **rmfbvol** command to delete the designated volume configurations if you do not want to obtain additional information about what caused the configuration error.

Note: In the majority of instances, this is the only method for correcting a configuration error.

Creating fixed block volume groups using the DS CLI

Complete this task to create fixed block volume groups.

A volume group identifies the set of fixed block logical volumes that are accessible by one or more SCSI host system ports. SCSI host system access is constrained to the identified access mode. Only those SCSI host ports that are registered to a volume group ID are allowed to access the set of logical volumes that is contained by the volume group.

Logical volumes can be assigned to a volume group when the volume group is created, or the logical volumes can be added (or removed) at a later time. The volume group type determines the maximum number of volumes that can be assigned to a volume group, either a maximum of 256 volumes or a maximum of 64 000 volumes. The volume group type must be selected according to the addressing capability of the SCSI host system that will use the volume group.

Perform the following steps to create and view fixed block volume groups:

1. Issue the **mkvolgrp** command to create a fixed block volume group. Enter the **mkvolgrp** command at the dscli command prompt with the following parameters and variables:

Note: Repeat this step for each volume group that you want to create.

```
dscli>mkvolgrp -dev IBM.1750-68FA120 -hosttype pSeries  
-volume 0001-0010,0120 my_nickname
```

Notes:

- a. You can use the **-hosttype** parameter with the **mkvolgrp** command. This parameter is an easier way of specifying the type of volume group. If you do not use the **-hosttype** parameter, it is assumed that the volume group type is **scsimask**.
 - b. You cannot use the **-type** parameter and the **-hosttype** parameter together.
 - c. If your volume group is not **scsimask** type and you do not want to use the **-hosttype** parameter, use the **-type** parameter. **scsimask** as the default value of the **-type** parameter; you can also specify **scsimap256** or **os400mask** as your volume group type. Because you need to know the criteria that is associated with these volume group types, see the **mkvolgrp** command for more information.
 - d. Volume IDs must meet the following criteria:
 - ID ranges must be separated by a comma (displayed as 0001-0010,0120 in the example).
 - For **scsimap256**, the array or ranges cannot exceed 256 volume ID entries. Otherwise, up to 64 384 entries are allowed.
 - Use the **-type 0s400mask** parameter if the volume group is limited to fixed block volume OS400-protected or OS400-unprotected types. Otherwise, the volume group is limited to the fixed block volume type 1750.
 - The volume group name (*my_nickname* in the example command) must be unique within the scope of the specified storage unit.
2. Issue the **lsvolgrp** command to create a list of assigned volume group IDs. Enter the **lsvolgrp** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsvolgrp -dev IBM.1750-68FA120 -l
```

Notes:

- a. The **lsvolgrp** command with the **-l** parameter displays a report with the following 3 values:
 - Name (the unique name that you assigned to the volume group)
 - Volume group ID
 - Type (the configured volume group type)

- b. You can narrow the scope of the report by requesting a specific type of volume. See the **lsvolgrp** command for information about the **-type** parameter.

Configuring fibre-channel I/O ports using the DS CLI

Complete this task to configure fibre-channel I/O ports using the DS command-line interface.

Before you begin, you must have the command-line interface prompt, and you must be connected to a storage unit that will be used for open systems host system storage.

In this process, you must designate the topology for the I/O port. The following three topology settings are available:

fibre channel arbitrated loop (coded as fc-al in the setiport command)

Enables the SCSI ULP with a FC-AL topology. The FC-AL topology does not support PPRC path I/O operations.

SCSI-FCP (coded as scsi-fcp in the setiport command)

Enables the SCSI ULP with a point-to-point or switched fabric topology. PPRC path I/O operations are enabled for this setting.

Note: Designate this topology for System i systems using i5/OS level V5R3M5 and above.

ficon (coded as ficon in the setiport command)

Enables the FICON ULP with a point-to-point or switched fabric topology. PPRC path I/O operations are not supported for FICON ULP.

The storage unit supports the fibre-channel host adapter (HA) card type. For machine type 1750, one or two HA cards are installed in each of the two CEC assemblies. Use the **lsiport** and **setiport** commands to configure fibre-channel I/O ports.

Each fibre-channel HA card contains four I/O ports. The storage image microcode automatically creates one I/O port to represent each HA card I/O port. The default fibre-channel I/O port settings enable SCSI-FCP “identified” access to fixed block volumes. You might have to modify the I/O port settings to enable SCSI FC-AL access to fixed block volumes.

To configure the fibre-channel ports, perform the following steps:

1. View a list of fibre-channel port IDs by typing the following command format at the dscli command prompt:

```
dscli>lsiport -dev IBM.1750-75FA120 -l -type fc
```

A detailed report is displayed that lists the fibre-channel I/O ports.

2. Analyze the report and determine which I/O port IDs that you want to access the fixed block volumes.

Configure a minimum of four I/O ports for SCSI host I/O operations. Select ports with physical locations on different host bus adapter (HA) cards. If possible, locate the HA cards in different I/O enclosures.

3. Set the I/O ports that you have identified to enable the FC-AL (fibre-channel arbitrated loop), SCSI-FCP, or FICON topology. The following example shows how to enable the FC-AL topology by typing the following command format at the dscli command prompt:

Note: I/O ports are automatically set to the offline state and returned to the online state after configuration changes are applied.

```
dscli>setioport -dev IBM.1750-75FA120 -topology fc-  
0012 0013 0112 0113
```

4. Press Enter. A successful process returns a confirmation message indicating that the port IDs have been successfully configured.

Creating SCSI host port connections using DS CLI

Complete this task to create SCSI host port connections using the DS command-line interface.

Before you begin, you must have the command-line interface prompt, and you must be connected to a storage unit that can be used for open systems host system storage.

The 1750 machine type supports the “identified” access mode for SCSI host attachments, which requires that all SCSI host ports be identified to a storage unit. One SCSI host port connection must be created for each SCSI host port that accesses storage unit fixed block volumes. Use the **lshosttype**, **mkhostconnect**, and **lshostconnect** commands to create the SCSI host port connections.

A SCSI host port contains attributes that identify the following information:

- SCSI host system type
- Port profile
- Port WWPN
- Volume group ID that the port accesses
- An array of storage unit I/O port IDs that the host port logs into for volume access
- An attribute to indicate that all I/O ports can be used for volume access
- Host port description
- Port nickname

There are two ways that you can approach this task:

- Use the **-hosttype** parameter with the **mkhostconnect** command. Using the **-hosttype** parameter is the best solution for most users.
- Use the **mkhostconnect** command with the **-lbs**, **-addrdiscovery**, and **-profile** parameters.

Notes:

1. Specifying the **-hosttype** parameter automatically sets the **-lbs**, **-addrdiscovery**, and **-profile** values.
2. If you do not use the **-hosttype** parameter, you must issue the **lspportprof** command to ensure that you obtain the correct values to use with the **-lbs**, **-addrdiscovery**, and **-profile** parameters.
3. You cannot use the **-hosttype** parameter with these other parameters.

The following task is described from the assumption that you have used the **-hosttype** parameter.

To configure the SCSI host ports, perform the following steps:

1. Obtain your host type information by issuing the **lshosttype** command. Enter the **lshosttype** command at the dscli command prompt with the following parameters and variables:

```
dscli>lshosttype -l -type volumeGroup_type
```

This command displays a report like the following:

Name	Profile	AddrDiscovery	LBS	Description
pSeries	IBM pSeries - AIX	reportlun	512	IBM pSeries, RS/6000 and RS/6000 SP Servers (AIX)
zLinux	IBM zSeries - zLinux	lunpolling	512	IBM zSeries Servers (Linux)
iSeries (if os400all was specified)	IBM iSeries - os400	reportlun	520	IBM iSeries Servers (System i)

Note: Volume group type is one of the following designations (use a separate command for each choice):

- ficonall
- scsiall
- scsimask
- scsimap256
- os400all
- os400mask

The same results are displayed when you specify os400all or os400mask or when you specify scsiall and scsimask or scsimap256.

2. Create SCSI host ports by issuing the **mkhostconnect** command. Enter the **mkhostconnect** command at the dscli command prompt with the following parameters and variables:

```
dscli>mkhostconnect -dev storage_image_ID -wwname wwpn
-hosttype host_type -volgrp volume_group_ID -ioport port_ID
host_name
```

Notes:

- a. The **-wwname** parameter specifies the 16-character worldwide name that is assigned to the host system fibre-channel adapter port. This WWPN value is validated each time that the host system port logs into an I/O port.
- b. The **-hosttype** parameter specifies fibre-channel communications layer characteristics that might be unique according to the host system manufacturer, operating system, or version of the system. Typical specifications are iSeries, pSeries, and so on.
- c. The **-volgrp** parameter specifies the volume group ID that this host port can access. Host port objects might be created prior to creating volume groups, in which case you must use the **chhostconnect** command to add volume group ID assignments at a later time.
- d. The **-ioport all** specifies SCSI host port (WWPN) access into all IO ports that are configured for the FC-AL or SCSI-FCP topology.
- e. **host_name** specifies the SCSI host system nickname that you have assigned.

3. Repeat Step 2 for each SCSI host system port that will access LUN volumes.
4. Verify that all SCSI host ports have been configured and that they are recognized by the storage unit according to your specifications by issuing the `lshostconnect` command with the `-I` parameter.

Configuring new count key data storage using DS CLI

This section describes how to configure new count key data (CKD) storage within your storage unit for a zSeries host system.

Before you begin, you must be logged into the DS CLI application in interactive command mode. You must also be connected to a storage unit that is used for zSeries systems host storage.

Configuring CKD storage involves two basic processes: the creation of the CKD storage configuration and the configuration of the storage unit I/O ports for zSeries host system attachment. These two basic processes can be performed in the reverse order, but it is better to create storage configurations first. Creating the storage configuration first creates the media to back up configuration data that is not related specifically to the storage configuration.

Configuring new CKD storage involves the following processes:

- Creating CKD extent pools
- Creating arrays
- Creating and associating ranks with extent pools
- Creating logical control units
- Creating CKD volumes
- Creating CKD volume groups (system generated).

The internal microcode automatically creates the CKD FICON All volume group ID (V10) and automatically assigns all CKD base and alias volumes to this volume group. This volume group ID (V10) is automatically assigned to storage unit I/O fibre channel ports that are configured for FICON I/O operations.

- Configuring fibre-channel I/O ports

Creating count key data extent pools using the DS CLI

Complete this task to create CKD volume extent pools. This is the first step in configuring new count key data storage. You can use the DS CLI commands to create extent pools for CKD volumes.

Creating the extent pools before the arrays and ranks saves a processing step. When you create the new ranks, you can assign them to existing extent pools. Otherwise, you must modify each rank object to complete the extent pool ID assignment after the extent pools have been defined.

Each extent pool is defined with the rank group of 0 or 1 and a storage type of **ckd**. At a minimum, you must define one extent pool for each rank group and storage type combination. This means that you must make a minimum of two extent pools for a storage unit that contains CKD storage: one CKD extent pool per rank group.

Extent pools that are defined for rank group 0 or 1 are assigned an even- or odd-numbered extent pool ID, respectively. Even-numbered extent pools are managed by storage server ID 0. Odd-numbered extent pools are managed by

storage server ID 1. Each rank group is assigned to one extent pool; therefore, storage server workload is affected by the rank assignments to even- and odd-numbered extent pool IDs. It is better to evenly distribute rank and extent pool allocations in order to keep the storage server workloads balanced.

Notes:

1. You can create more than the minimum number of extent pools. For example, you can define unique extent pools for each RAID type (5 or 10) that is configured in a storage image. Or, you can define and name extent pools according to the host system attachments that access the volumes that are created from extent pool extents.
2. You can have the same number of extent pools as ranks.

Use the **lsextpool** and **mkextpool** commands to create the CKD extent pools. You must be logged into the DS CLI application and connected to the storage unit that will be used for your zSeries host system.

To make your extent pools, perform the following steps:

1. **Find defined CKD extent pools** by issuing the **lsextpool** command to display a list of the existing CKD extent pools. Type the following command at the dscli command prompt:

```
dscli> lsextpool -dev IBM.1750-75FA120 -stgtype ckd
```

2. **Analyze the extent pool listing** for the following information.

- Does the minimum set of extent pools exist? There must be one extent pool for rank group 0 and one extent pool for rank group 1.

Note: If this is the first time that the extent pools are created, the minimum number of extent pools does not already exist.

- Does each extent pool have a rank group that is assigned to it and are they balanced? If this is the first time, there will be no rank assignments.
- Are additional extent pools needed?

3. **Make your extent pools** by issuing the **mkextpool** command. (A minimum of two extent pools must be created. One for rank group 0 and one for rank group 1.)

Type the following command at the dscli command prompt:

```
dscli>mkextpool -dev IBM.1750-68FA120 -rankgrp 0 -stgtype ckd  
extent_pool_name  
dscli>mkextpool -dev IBM.1750-68FA120 -rankgrp 1 -stgtype ckd  
extent_pool_name
```

The *extent_pool_name* parameter is required with the **mkextpool** command. The extent pool name is a unique name that you assign to each extent pool and it cannot be longer than 16 characters.

Make additional extent pools for each of the following conditions:

- Each RAID type (5 or 10)
- Each disk drive module (DDM) size
- Each CKD volume type (3380, 3390)
- Each logical control unit (LCU) address group

4. **Press Enter.** A successful process displays the following message:

Note: The unique name that you assigned to the extent pool does not display in the process message. However, when you issue the **lsextpool** command, the extent pool name is displayed.

Extent pool P1 successfully created.

5. **Repeat Step 2** for each extent pool that you want to create. Try to evenly distribute rank and extent pool allocations in order to keep the storage server workloads balanced.
6. **Verify the extent pool assignments** by issuing the **lsextpool** command when you are done making the extent pools. Use the **-l** parameter to display a full report for the extent pools that are assigned to the storage unit. Type the following command at the dscli command prompt:

```
dscli> lsextpool -dev IBM.1750-75FA120 -l
```

You might want to print this report because this list is used later during the process of creating CKD volumes.

Creating arrays for CKD volumes using the DS CLI

Complete this task to create arrays for CKD volumes using the DS CLI commands.

The machine type 1750 must contain at least one storage enclosure, with a minimum of four DDMs.

The DDMs of a storage enclosure are partitioned into array sites. A machine type 1750 array site consists of four DDMs in one storage enclosure of a storage enclosure pair, with two to eight (four DDM) array sites per storage enclosure pair. All array sites of a storage enclosure pair have identical capacity, rpm, and interface characteristics, and an interface to a common DA pair.

The creation of arrays is based on the array sites that are associated with the storage unit. You must make an array from 1 or 2 array sites. An array inherits the characteristics of its parent array sites, and is given a RAID type attribute (5 or 10). A 1750 Array object of RAID type 5 or 10 is made from one or two (4 DDMs) array sites.

Note: The array status is “unassigned” until the array is assigned to a rank.

Use the **lsarraysite** and **mkarray** commands to create the arrays. You must be logged into the DS CLI application and connected to the storage unit that will be used for open systems host system storage.

Perform the following steps to create arrays for a CKD volume configuration:

1. Issue the **lsarraysite** command to find the unassigned array sites. Enter the **lsarraysite** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsarraysite -dev storage_image_ID -state unassigned
```

Note: If this is your first time creating volumes, you will see all the arrays with a state of “unassigned”.

2. Press Enter. A report of unassigned array sites is displayed. Use the list to identify unassigned array site capacity, rpm, and device adapter (DA) pair attributes. Record the RAID type for each array site.
3. Issue the **mkarray** command to create an array from each site with the status “unassigned”. Enter the **mkarray** command at the dscli command prompt with the following parameters and variables:

```
dscli>mkarray -dev storage_image_ID -raidtype [5 | 10] -arsite array_site
```

Repeat this command until all unassigned array sites have been assigned to an array.

Notes:

- a. You can specify one or two array sites for RAID types 5 and 10. If there are two array sites, both sites must be associated with a common DA pair ID. Two array sites must be separated by commas with no blank space in between. Example: S10,S11.
- b. The new array site inherits the capacity, rpm, interface, and DA pair characteristics of its parent array sites. The state of the array is “unassigned” until it is assigned to a rank.

Creating a rank for CKD volumes using the DS CLI

Complete this task to create a rank for a CKD volume. A rank is a logically contiguous storage space that is made up of one or more arrays. You want to assign a rank to every unassigned array.

A rank inherits the characteristics, including RAID type, of its parent array and is given a storage type attribute FB (fixed block) or CKD (count key data). The rank configuration state is unassigned until it is assigned to an extent pool. An unassigned rank is not associated with either rank group 0 or 1. Any unassigned rank can be assigned to an extent pool that is associated with either rank group 0 or 1.

Note: You can assign a rank to an unassigned array and also assign the rank to an extent pool at the same time if you have already created the extent pools and the arrays. Creating extent pools first saves a step in the configuration.

Use the **lsarray**, **mkrank**, and **lsrank** commands to assign a rank to each unassigned array. You must be logged into the DS CLI application and be connected to the storage unit that will be used for open systems host system storage.

To make ranks, perform the following steps:

1. Issue the **lsarray** command to ensure you have a list of the unassigned arrays for which ranks must be assigned. Enter the **lsarray** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsarray -dev IBM.1750-68FA120 -state unassigned
```

2. Issue the **mkrank** command to assign a rank to rank group 0 or 1 according to the rank group number of the assigned extent pool ID. Enter the **mkrank** command at the dscli command prompt with the following parameters and variables:

```
dscli>mkrank -dev IBM.1750-68FA120 -array A44  
-stgtype ckd -extpool P1
```

Notes:

- a. You can specify either the **-wait** or the **-extpool** parameter when you use the **mkrank** command. Either of these parameters allows you to be notified if the rank configuration has failed for any reason.
 - b. Be aware that when you use the **-wait** parameter, you cannot issue any other commands until the entire transaction has processed.
3. Press Enter to create the ranks.

The process of making the rank involves formatting drives. It can take a little time before the process finishes. To check on the process, issue the **lsrank** command from a different DS CLI session. A successful process generates the following type of message:

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Device: IBM.1750-68FA120
```

```
Rank IBM.1750-75FA120/R44 successfully created.
```

4. Repeat Step 2 and step 3 until all unassigned arrays are assigned a rank and an extent pool.
5. Issue the **lsrank** command to verify that ranks and extent pools have been assigned. Enter the **lsrank** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsrank -dev IBM.1750-68FA120 -l
```
6. Press Enter. A report of the rank assignments for your entire storage unit is displayed.

Creating logical control units for CKD volumes using DS CLI

The logical control unit (LCU) is the S/390 and zSeries host equivalent of the logical subsystem (LSS) for open systems hosts. The LCU must be defined (created) before CKD logical volumes can be created.

The DS6000 has a 16 384 volume address space that is partitioned into 64 logical subsystem (LSS) units, where each LSS contains 256 logical volume numbers. The 64 LSS units are assigned to one of 4 address groups, where each address group contains 16 LSSs, or 4096 volume addresses. All of the LSSs in one address group must be of the same type (CKD or fixed block).

Typically, LCUs are created in groups of 16, beginning at LSS address X'x0'.

Use the **lsaddressgrp**, **mklcu**, and **lslcu** commands to create the LCU type logical subsystems. You must be logged into the DS CLI application and connected to the storage unit that will be used for open systems host system storage.

To create LCUs, perform the following steps:

1. Find unassigned and available address groups by issuing the **lsaddressgrp** command. To use the **lsaddressgrp** command, type the following at the dscli command prompt:

```
dscli>lsaddressgrp -dev IBM.1750-75FA120
```

This command displays a report on the status of the address groups within your storage unit.
2. Analyze the report to identify all of the address groups that are available to be defined. Use the following criteria:
 - If the list is empty, all of the address groups are available to be defined.
 - A defined address group with the storage type fb (fixed block) is not available to be defined.
 - A defined address group with the storage type ckd and with fewer than 16 LSSs is available for LCU definition.
 - If you are using an undefined address group to make new LCUs, select the lowest numbered address group that is not defined.
 - If you are defining a new LCU in an existing CKD address group, use the **lslcu** command to identify LCUs that are already defined in the target address group.

3. Make the LCU logical subsystem objects by issuing the **mklcu** command. Type the command using the following format at the dscli command prompt:

```
dscli>mklcu -dev IBM.1750-75FA120 -qty 16 -id 00 -ss 0010 -lcutype 3390-3
```

In this example, the values specify the following:

qty

Specifies the number of LCU IDs to be created.

id Specifies the LCU ID to be created, or the first LCU ID in a sequence of LCU IDs to be created.

ss Specifies the subsystem ID that you have assigned. If multiple LCU IDs are being created, then the SSID value increments for each additional LCU ID that is created.

If 16 LCUs are created, starting with SSID 0x10, then the SSID values are 0x0010 – 0x001F.

lcutype

Specifies the type of LCU to be created. You can specify the following types:

- 3390-3
- 3990-tp
- 3990-6
- bs2000

4. Press Enter. A successful process displays a confirmation message listing each LCU ID number that has been successfully created.
5. Verify that the LCUs are recognized in the storage unit by issuing the **lslcu** command at the dscli command prompt as follows:

```
dscli>lslcu -dev IBM.1750-75FA120 -l
```

Using the **-l** parameter displays a more detailed report for each LCU that is associated with your storage unit.

Creating count key data volumes using the DS CLI

Complete this task to create count key data (CKD) volumes.

A logical volume consists of one or more data extents that are allocated from a single extent pool. The volume data type is inherited from the extent pool extent storage type (fixed block or CKD) characteristic. When a CKD volume is created, volume attributes are further defined by a base or alias volume type, 3390 or 3380 volume cylinder type, and volume capacity in cylinders. These volume attributes characterize the volume to the host system that will eventually access the volume. Each volume is assigned a volume ID, which is the volume address within the 64 KB address space. Host access to a volume is enabled when the volume ID is assigned to a volume group; however, CKD volumes are automatically assigned to the volume group CKD FICON/ESCON All (ID V10).

Perform the following steps to create your CKD volumes:

1. View your list of CKD extent pool IDs and determine which extent pool IDs that you want to use as the source for the CKD volumes to be created. You obtained this list when you first created your extent pools. If this list is not available, you can issue the **lsextpool** command to obtain the list of extent pool IDs.

Extent pool attributes determine the size and quantity of volumes that can be created. The extent pool ID (even/odd) indicates the storage server (0|1), which dictates that the logical control unit (LCU) ID component of the volume ID must be an even or an odd number.

2. Issue the **mkckdvol** command to make 128 base volumes for each LCU. Enter the **mkckdvol** command at the dscli command prompt with the following parameters and variables:

```
dscli>mkckdvol -dev IBM.1750-68FA120 -extpool p1 -cap 3339  
-name finance#d 0000-007F
```

The following considerations affect the command example in this step:

- The **-extpool** parameter identifies a CKD extent pool that contains available data extents.
- The **-cap** parameter specifies the quantity of CKD cylinders that are allocated to this volume.
- The **-name** parameter allows you to assign an easy-to-use label or nickname to the volume. The volume name parameter can include a wild card (#d or #h) that inserts a decimal or hexadecimal volume ID value into the volume name.

Note: The decimal designation does not apply to the volume ID number or the number of volumes that were created by the command. It only applies to the unique name that you have assigned to the volume. When you process the **mkckdvol** command, the volume name that you have assigned does not appear in the confirmation message. To view the volume name that you have assigned, issue the **lsckdvol** or **showckdvol** command.

- Volume ID 0000 - 007F specifies 128 volumes, starting at CKD address group (0), LCU ID (00), and volume number (00). You must specify volume IDs that have not been previously defined as CKD or fixed block volumes.
3. Press Enter to create the volumes. A confirmation message is displayed that lists the successful creation of each volume.
 4. Repeat Steps 2 and 3 until all required logical volumes for all LCUs have been created.

5. Issue the **mkaliasvol** command to make 128 alias volumes for each LCU. Enter the **mkaliasvol** command at the dscli command prompt with the following parameters and variables:

```
dscli>mkaliasvol -dev IBM.1750-68FA120 -base 0000-004F  
-order decrement -qty 2 00FF
```

Consider the following conditions with regard to the command example in this step:

- The **-base 0000 - 004F** parameter specifies that alias volumes are assigned to existing base volume IDs 0000 - 004F. Base and alias volumes must be associated with a common LCU ID.
- The **-order** parameter specifies the order in which alias volume IDs are assigned.
- The **-qty** parameter specifies the number of alias volumes that are assigned to each base volume.
- The **volume ID (00FF)** parameter specifies that the alias volumes are assigned, starting at a CKD address group (0), LCU ID (00) and volume number (FF). You are responsible for specifying the volume ID values that have not been previously defined as CKD or fixed block volume types.

As a result, alias volumes 00FF and 00FE are created for base volume 0000, 00FD and 00FC for 0001, and so on.

6. Repeat Step 5 until you have defined all required logical volumes for all the LCUs.
7. Press Enter to create the alias volumes. A confirmation message is displayed that lists the successful creation of each volume.
8. Issue the **lsckdvol** command to display a report that you can use to confirm the status of your CKD volumes. Enter the **lsckdvol** command at the dscli command prompt with the following parameters and variables:
dscli>lsckdvol -dev IBM.1750-6800861 -l 1410

Note: It is possible that the report will display that there was a configuration error that is associated with one or more of your **mkckdvol** transactions. In the majority of instances, the only way to correct this error is to issue the **rmckdvol** command.

Correcting a CKD volume configuration error

Complete this task to correct a count key data (CKD) volume configuration error.

There might be occasions when you are using the **mkckdvol** command to create CKD volumes, but the transaction fails. You might not be aware of the failure until you run the **lsckdvol** or the **showckdvol** command to check the status of the volumes that you have created.

The **lsckdvol** or the **showckdvol** commands display reports that includes a *configstate* category. The configuration state category reports on the current state of the rank. One of the configuration state codes is **configuration error**.

A status of **configuration error** specifies that the configuration process did not complete successfully. This state reflects an internal error condition and is not an indication that there was a user input error.

You might want to gather additional information about what caused the error, which can help you determine how to correct it. Generally, to correct this error state, you must delete the designated volume configuration and submit a new transaction request.

Perform the following steps to obtain additional information about the configuration error and to correct this error condition.

1. Add the **-v** (verbose) command flag to your **mkckdvol** command and reissue the **mkckdvol** command for the transactions that show there is a configuration error.

Note: You can also turn on the verbose mode in your profile file and reissue the command.

If you designate the verbose mode, the display of extra output includes the error code that is generated when the create CKD volume transaction fails.

2. Issue the **rmckdvol** command to delete the designated volume configurations if you do not want to obtain additional information about what caused the configuration error.

Note: In the majority of instances, this is the only method for correcting a configuration error.

Configuring fibre-channel I/O ports using the DS CLI

Complete this task to configure fibre-channel I/O ports using the DS command-line interface.

Before you begin, you must have the command-line interface prompt, and you must be connected to a storage unit that will be used for open systems host system storage.

In this process, you must designate the topology for the I/O port. The following three topology settings are available:

fibre channel arbitrated loop (coded as fc-al in the setioport command)

Enables the SCSI ULP with a FC-AL topology. The FC-AL topology does not support PPRC path I/O operations.

SCSI-FCP (coded as scsi-fcp in the setioport command)

Enables the SCSI ULP with a point-to-point or switched fabric topology. PPRC path I/O operations are enabled for this setting.

Note: Designate this topology for System i systems using i5/OS level V5R3M5 and above.

ficon (coded as ficon in the setioport command)

Enables the FICON ULP with a point-to-point or switched fabric topology. PPRC path I/O operations are not supported for FICON ULP.

The storage unit supports the fibre-channel host adapter (HA) card type. For machine type 1750, one or two HA cards are installed in each of the two CEC assemblies. Use the **lsioport** and **setioport** commands to configure fibre-channel I/O ports.

Each fibre-channel HA card contains four I/O ports. The storage image microcode automatically creates one I/O port to represent each HA card I/O port. The default fibre-channel I/O port settings enable SCSI-FCP “identified” access to fixed block volumes. You might have to modify the I/O port settings to enable SCSI FC-AL access to fixed block volumes.

To configure the fibre-channel ports, perform the following steps:

1. View a list of fibre-channel port IDs by typing the following command format at the dscli command prompt:

```
dscli>lsioport -dev IBM.1750-75FA120 -l -type fc
```

A detailed report is displayed that lists the fibre-channel I/O ports.

2. Analyze the report and determine which I/O port IDs that you want to access the fixed block volumes.

Configure a minimum of four I/O ports for SCSI host I/O operations. Select ports with physical locations on different host bus adapter (HA) cards. If possible, locate the HA cards in different I/O enclosures.

3. Set the I/O ports that you have identified to enable the FC-AL (fibre-channel arbitrated loop), SCSI-FCP, or FICON topology. The following example shows how to enable the FC-AL topology by typing the following command format at the dscli command prompt:

Note: I/O ports are automatically set to the offline state and returned to the online state after configuration changes are applied.

```
dscli>setioport -dev IBM.1750-75FA120 -topology fc-al  
0012 0013 0112 0113
```

- |
4. Press Enter. A successful process returns a confirmation message indicating that the port IDs have been successfully configured.

Chapter 9. Upgrading the DS6000 code level

Complete this group of tasks to locate, download, and install code upgrades for your storage unit, DS Storage Manager, and DS CLI. You can perform either a concurrent (DS6000 has current I/O activity) or nonconcurrent (no current I/O activity on the DS6000) code load.

If you are performing a concurrent code upgrade, you must first ensure that all host paths are available and operating correctly. During a concurrent code upgrade, the storage unit host ports temporarily lose their connections for each processor card when that card is upgraded. Before you begin a concurrent code upgrade, ensure that each host has a path to each processor card to prevent a lost connection between the host and the storage unit. If the host does not have a path to each processor card, the paths from the host to the storage unit are lost during the code upgrade.

Note: While each host port is temporarily unavailable during the code upgrade, I/O operations from a host can cause errors to be logged stating that the path is lost. Error message activity can make the host interface unavailable for other host system administration tasks during the code upgrade.

The DS6000 operates its various components with upgradeable code (or firmware). This code governs the operation of the various parts of the unit, such as the device adapters, host adapters, processor cards, and other advanced function features. As IBM continues to develop and improve this code, code updates will be available to you.

Perform the following steps for each storage unit that is connected to the DS Storage Manager.

1. Check for firmware updates. Follow the instructions in “Checking for code updates (real-time only)” on page 200. If you are not required to upgrade your code, stop here.
2. To upgrade the code successfully, before you proceed with the code upgrade you must ensure that no alert LEDs are illuminated and that there are no open problems in the problem logs. Correct any problems before you proceed with the firmware upgrade. You can correct problems by performing the steps in either “Following a light path to perform unguided service” on page 201 or “Performing guided service through the problem log” on page 202.
3. Obtain the downloadable files from the DS6000 support Web site. Follow the instructions in “Downloading code updates” on page 203.
4. Install the code update on your storage unit. Follow the instructions in “Installing code upgrades (real-time only)” on page 203. You must perform this step for each storage unit that is managed by your DS Storage Manager before you proceed to the next step.
5. Upgrade the DS Storage Manager. Begin the installation of the new DS Storage Manager by completing one of the following sets of steps:
 - If you are installing from the ZIP file bundle, follow these steps:
 - a. Extract the DS Storage Manager ZIP folder that is located in the file path where you extracted the ZIP file bundle information.
 - b. Navigate to the folder that you just extracted and run setup.exe. You can perform the upgrade using either the graphical mode or silent mode. See

- “Upgrading the DS Storage Manager on a Windows operating system using the graphical mode” on page 204 or “Upgrading the DS Storage Manager on the Windows operating system in unattended (silent) mode” on page 208 for the steps. When the installation is complete, a prompt to restart might be displayed. If possible, do not restart at this time. You will perform a restart after you install the DS CLI upgrade.
- If you are installing from the ISO image that you used to create an installation CD, insert the CD and navigate to the DS Storage Manager installation files and run setup.exe. You can perform the upgrade using either the graphical mode or silent mode. See “Upgrading the DS Storage Manager on a Windows operating system using the graphical mode” on page 204 or “Upgrading the DS Storage Manager on the Windows operating system in unattended (silent) mode” on page 208 for the steps. When the installation is complete, a prompt to restart might be displayed. If possible, do not restart at this time. You will perform a restart after you install the DS CLI upgrade.
6. Perform the following steps to upgrade the DS CLI:
 - a. Begin the installation of the new DS CLI by completing one of the following sets of steps:
 - If you are installing from the ZIP file bundle, follow these steps:
 - 1) Extract the DS CLI ZIP folder that is located in the file path where you extracted the ZIP file bundle information.
 - 2) Navigate to the folder that you just extracted. To install using graphical mode, run the file setupwin32.exe. To install using console mode, run the file setupwin32console.exe.
 - If you are installing from the ISO image that you used to create an installation CD, insert the CD and navigate to the DS CLI installation files. To install using graphical mode, run the file setupwin32.exe. To install using console mode, run the file setupwin32console.exe.
 - b. Follow the instructions for upgrading the DS CLI. See “Installing the DS CLI using the graphical mode” on page 119 or “Installing the DS CLI using the console mode” on page 113 to perform the upgrade using either the graphical mode or console mode.
 7. If you have not already done so, restart your management console.
 8. After the DS Storage Manager is installed, use Storage units — Main page to verify that you can view the status of the storage complex and its associated logical configuration. In the navigation, select **Real-time manager** → **Manage hardware** → **Storage units**. On the Storage units - Main Page, select the appropriate storage unit.
 9. Ensure that no alert LEDs are illuminated and that there are no open problems in the problem logs. Correct any problems. You can correct problems by performing the steps in either “Following a light path to perform unguided service” on page 201 or “Performing guided service through the problem log” on page 202.

If you are setting up the DS Storage Manager, return to “Setting up the DS Storage Manager” on page 90 and complete the rest of the steps.

Checking for code updates (real-time only)

Complete this task to view your current code level and determine if you need to upgrade to a newer level of code.

Perform the following steps to compare the firmware level that is currently installed on your storage unit with the available firmware updates to determine if you need to upgrade the code:

1. In the navigation, select **Real-time manager** → **Manage hardware** → **Storage units**. On the Storage units - Main Page, select the appropriate storage complex, then select the appropriate storage unit.

Note: You can select only one storage unit at a time. To upgrade more than one storage unit, you must complete this task for each storage unit.

2. From the **Select Action** list, select **Apply firmware update**. The Apply firmware update page is displayed.
3. Locate the Current firmware level row in the table to view your current level of code.
4. Click **Check for firmware updates**. The Check for firmware updates page is displayed.
5. Click the link that is provided on the Check for firmware updates page. The DS6000 support Web site is displayed.
6. View the downloadable files to find the highest level of code that is available.
7. Compare your current level of code with the available levels of code to determine if you must perform a code upgrade. You must upgrade your code if the available level of code is higher than your current level.

Following a light path to perform unguided service

Complete this task to locate and repair a resource using the LED indicators that are located throughout the enclosure.

You can use these steps to follow LED indicators from the beginning to the end of a service action. You do not need to access the DS Storage Manager to replace a resource that has a solid amber error light.

1. Find the amber system alert light on the front display panel of the server enclosure.
2. Determine where the fault is located and perform one of the following actions:
 - If the fault on opposite side indicator is lit on the front display panel, move to the rear of the server.
 - If the fault in external enclosure indicator is lit on the front display panel, use the location indicators and the **Identify** button on the rear operator panel to find the enclosure that is connected to the server with a solid amber system alert light.
 - If the system alert indicator is lit and no other lights or indicators are lit on the front display panel, the fault is on a resource that is found on the front of the enclosure. Continue to step 3.
3. Find the resource with a solid amber error indicator.
4. Order a new resource to replace the one with the fault.
5. Remove the resource.

Attention: Review the remove and replace procedures for a resource before removing it. You must replace some resources within a specific time limit. In some cases, if you remove a resource without immediately replacing it, you must place a blank version of the resource in the empty slot to prevent overheating of the storage unit.

6. Replace the resource. The system automatically senses the new resource and begins the procedure to resume normal operation for that resource.

7. Determine if you need to perform further repair actions.
 - If the system alert indicator is no longer lit, you have successfully replaced all required resources. The error entry in the system log closes and the storage unit resumes normal operation. You do not need to take any further action.
 - If, after you replace the resource, the system alert indicator is still lit, you must replace another resource. Repeat these steps to identify and replace the next resource.

Performing guided service through the problem log

Complete this task to locate and repair a resource using the log entries in the DS Storage Manager.

Before you begin, you must review the effects of performing a service action on the resource you are going to repair.

Perform the following steps to service a storage unit using the information found in the Logs pages of the DS Storage Manager.

1. In the navigation, select **Real-time manager** → **Monitor system** → **Logs**.
2. The initial log entry table displays new error or informational entries from the preceding 25 hours. If you want to view an entry for an event that occurred longer than 25 hours ago, you must select the appropriate storage unit information to filter the table.
3. **Select** the error entry for which you want to complete service.
4. Select **View details, Go**, to view more specific log entry information.
5. Find and select the first resource in the list of affected resources. If more than one resource must be replaced, the resources are listed in the order in which they must be replaced. Replace the first resource in the list before the others. Next, replace the second resource in the list. Follow this method until all required resources have been replaced.
6. Order a replacement resource, if necessary.
7. Select **Take resource offline** to begin to quiesce the resource.

Attention: Review the remove and replace procedures for a resource before quiescing and removing it. You must replace some resources within a specific time limit. In some cases, if you remove a resource without immediately replacing it, you must place a blank version of the resource in the empty slot to prevent overheating of the storage unit.
8. Find the appropriate resource in the storage unit.
9. Wait for the solid amber alert light to come on.
10. Remove the resource.
11. Replace the resource. The resource will automatically begin the procedure to resume operation.
12. Review the Logs page to determine if the log entry has closed or if you need to perform additional actions.

Note: If you take a hardware resource offline and decide not to remove it from the enclosure, you must bring the resource online through the DS Storage Manager before it can resume operation.

Downloading code updates

Complete this task to download the most recent code upgrade file as either a ZIP file bundle or an ISO CD-ROM image.

You must have an IBM user ID and password before you can download code updates. If you do not have a user account, you can create one at <http://www.storage.ibm.com/>.

To download code updates, perform the following steps:

1. From the Downloadable files section of the DS6000 support Web site, click the link for the highest level of code that is currently available. The DS6000 Microcode Release page for the selected code level is displayed.
2. On the DS6000 Microcode Release page, check the concurrent code load support chart to verify that your host types are supported.
3. Click the **Download** link to access the Code Download page. You must log in using your IBM user ID and password.
4. Select the version that you want to download, and click **Continue**.
5. Provide the required information, such as the target install date, the current code level, and the serial numbers of the storage units on which you are installing the code update, and then click **I agree**.
6. Read the release notes before you upgrade the code. Click **Download Now** to access the release notes.
7. Perform one of the following two sets of steps to download the code:
 - If you are installing the code update on a storage unit which is managed by the DS Storage Manager on which you are downloading code, perform the following steps:
 - a. Click **Download now** for the ZIP File Bundle. You can download through either http or the Download Director.
 - b. Save the ZIP bundle file.
 - c. Extract the compressed files from the ZIP file bundle. Make a note of where you save the extracted files, because you will have to provide that location when you install the update.
 - If you are downloading the code to create an installation CD that can be used for multiple upgrades, perform the following steps:
 - a. Click **Download now** for the ISO CD-ROM Image Bundle. You can download through either http or the Download Director.
 - b. Save the file. Make a note of where you save the file, because you will have to provide that location when you install the update.
 - c. Use the ISO CD-ROM image to create an installation CD.

Installing code upgrades (real-time only)

Complete this task to install upgraded code on your storage unit.

You must first download the code. You can download the code on your management console, or you can create an installation CD.

Perform the following steps to install code upgrades on the storage unit using the DS Storage Manager.

Important: You must upgrade all storage units that are managed by the same management console.

1. In the navigation, select **Real-time manager** → **Manage hardware** → **Storage units**. On the Storage units - Main Page, select the appropriate storage unit.
2. From the **Select Action** list, select **Apply firmware update**, and then click **Go**. The Apply firmware update page is displayed.
3. From the **Select a firmware application method** list, select either Concurrent or Non-Concurrent for the load method.

Note: For a noncurrent code load, ensure that there is no current I/O activity on the DS6000 before you proceed.

4. Under **Select firmware file to transfer to storage unit**, click **Select a file** to specify the code upgrade file. To search for the file, click **Browse**.
5. Locate and select the code upgrade file, and then click **OK**.
 - If you are installing from the ZIP file bundle, browse the files that you extracted from the bundle and select the SEA.jar file.
 - If you are installing from the ISO image that you used to create an installation CD, browse the files on the CD and select the SEA.jar file.
6. Click **Transfer file**. The system performs a compatibility check. This process can require a significant amount of time, depending on your network. When the process is complete, a message is displayed that either confirms a successful transfer or notifies you that the transfer has failed. If the transfer fails, verify that you are pointing to the correct file and try again to transfer the file.
7. When the file transfer is completed successfully, click **Activate** to initiate the code upgrade process. One of the following results occurs:
 - For a concurrent firmware update, the firmware update process is activated.
 - For a nonconcurrent firmware update, the Activate nonconcurrent firmware update page is displayed. Click **Proceed** to activate the firmware update process.

Note: The firmware update process can take 20 minutes to 2 hours, depending on your configuration.

Important: After you upgrade to the latest code level on all storage units, you must install the new DS Storage Manager and DS CLI software on the management console from the ISO image that you copied to the CD.

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:

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

1. Log on as a user with administrator authority.
2. 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:

DS6000 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.

DS6000 Storage Manager Web site

Provides information from the product Web site.

Installation wizard

Starts the IBM System Storage DS Storage Manager installation program.

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 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).

11. Click **Finish** to exit the installation process.

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

12. 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\DS6000 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 DS6000 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.

13. End the LaunchPad program by clicking **Exit** on the LaunchPad window. Restart if instructed to do so.
14. You have now completed the software upgrade portion of the DS6000. Continue to the post installation tasks.

Upgrading the DS Storage Manager on the Windows operating system in 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.

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.

Chapter 10. Removing the DS Storage Manager and DS CLI

This section contains the steps that you must complete to successfully remove either interface.

You can remove the DS Storage Manager and DS CLI using the same modes that are allowed by the operating systems during the installation process. For example, you can install the DS Storage Manager on a Windows system using the graphical mode. Conversely, you can remove the DS Storage Manager using the graphical mode.

Removing the DS Storage Manager

You can uninstall the DS Storage Manager using either the graphical mode or the unattended (silent) mode.

If you uninstall the DS Storage Manager in graphical mode, you are presented with a series of windows that guide you through the uninstallation process. If you uninstall the DS Storage Manager in unattended (silent) mode, you can use the command prompt to remove all installed files and restart your management console.

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 DS6000 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 V5 - 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.
9. 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.

1. Log on to your Windows operating system with administrator authority.
2. Open the Command Prompt window.
3. Navigate to the directory where the DS6000 Storage Manager is installed.
4. Type `<dest_path>_uninst\uninstall.exe -silent`, where `<dest_path>` is the path where the DS Storage Manager is installed.
5. Wait for the program to remove the DS Storage Manager.
6. 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.

Uninstalling the DS CLI

This section contains the tasks that you can follow to remove the DS CLI from your management console.

You can use one of the following methods, to remove the DS CLI from an open system:

- Silent mode
- Graphical mode
- Console mode

If you are removing the DS CLI from an OpenVMS system, you must use the console mode.

If you are removing the DS CLI from an OS/400 or iSeries system, you have the option to use the direct or remote method.

Removing the DS CLI using unattended (silent) mode

Use the unattended (silent) mode to remove the DS CLI through the command line if the DS CLI is installed on a UNIX system or a variant of UNIX (for example, HP-UX, Sun, or AIX).

Perform the following steps to successfully uninstall the DS CLI.

Notes:

1. If you are using Windows or Novell, you will use the Add/Remove Programs feature to uninstall the DS CLI.
2. This uninstall process only works with DS CLI. No other versions of CLI can be removed with this process.
1. Locate the uninstaller file in the `/_uninst` folder. If you selected the default directory, you can find the `_uninst` folder using the `/opt/ibm/dscli` path. The uninstaller file name is `uninstaller.xxx`, where `xxx` depends on the operating system. If you have a Hewlett Packard, Sun, or AIX system, then the file name is `uninstaller.bin`. For all other operating systems the file name is `uninstaller.sh`.
2. Type the following command at the command prompt: `<install directory>/_uninst/uninstaller.<exe | sh | bin> -silent`
3. Press the **Enter** key. All the associated CLI files are uninstalled.

Removing the DS CLI using the console mode

Use the console mode to remove the DS CLI when the DS CLI is installed on a UNIX system that does not have use of an X display.

Perform the following steps to remove the DS CLI using the console mode:

Note: Do not use the console method to uninstall DS CLI on a Windows system. Instead, follow the steps in this guide for removing the DS CLI using graphical mode.

1. Type the following command at a command prompt: `<install directory>/_uninst/uninstaller.<sh | bin> -console`
2. The Welcome screen displays. Press 1 and Enter to continue, or 3 to Cancel the removal process.

```
Welcome to the InstallShield Wizard for IBM System Storage DS Command-Line Interface (CLI)
The InstallShield Wizard uninstalls IBM System Storage DS Command-Line Interface on your computer.
```

```
To continue, choose Next.
```

```
DS Command-Line Interface
IBM Corporation
```

```
Press 1 for Next, 3 to Cancel or 4 to Redisplay [1]
```

3. The Uninstallation Location screen is displayed. Press 1 and Enter to continue, or 3 and Cancel to exit the removal process.

```
IBM System Storage DS Command-Line Interface will be uninstalled from the following location:
```

```
C:\Program Files\ibm\dsccli
```

```
Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]
```

4. The Uninstallation progress screen is displayed while the command-line interface is being removed.

```
Uninstalling IBM System Storage DS Command-Line Interface...
```

5. The Uninstallation Finish screen is displayed. Press 3 to finish the removal.

```
The InstallShield Wizard has successfully uninstalled IBM System Storage DS Command-Line Interface. Choose Finish to exit the wizard.
```

```
Press 3 to Finish or 4 to Redisplay [3]
```

Removing the DS CLI from your system using graphical mode

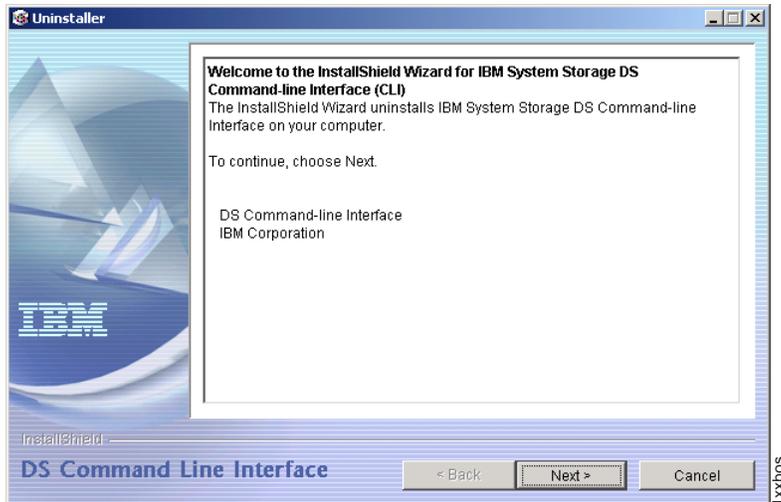
Use the graphical mode to remove the DS CLI from your system when the DS CLI is installed on a Windows, Novell, or UNIX system.

Notes:

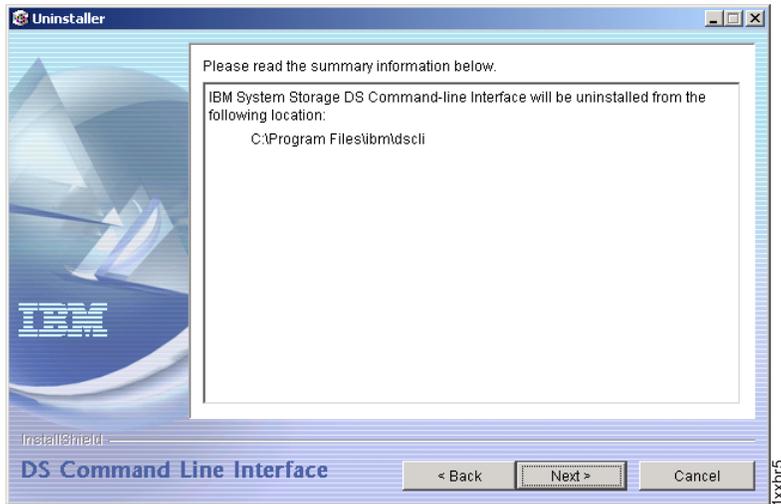
1. The following procedure applies to the removal of only the DS CLI application. This procedure cannot be used to remove other versions of the CLI.
2. If you do not want to create a new profile when you reinstall the CLI, select to not delete the DS CLI profile as you complete this task.

You use the Add/Remove Programs facility of the Windows operating system to remove the DS CLI from your system. When you have processed the uninstall steps, restart your system to complete the uninstall. Perform the following steps to remove the DS CLI using the graphical mode.

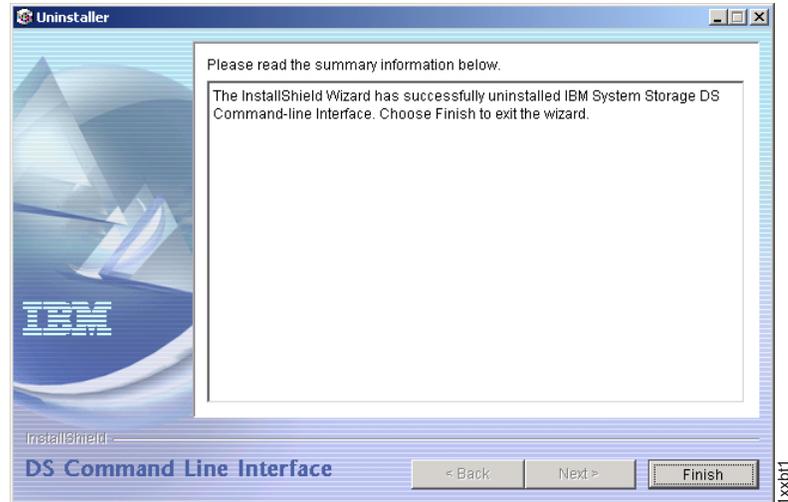
1. Navigate to the Windows Control Panel and open the Add/Remove program facility.
2. Scroll the list of currently installed programs and click the listing for IBM System Storage DS Command-Line Interface.
3. Click the **Change/Remove** button and the Welcome window for the Uninstaller is displayed.



4. 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 DS CLI program is removed.



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



If the uninstall program does not remove some information from your system, the Restart window is displayed. You must restart so that previously locked files are released and automatically deleted.

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

Removing the DS CLI from an OpenVMS system

The removal of the DS CLI application from an OpenVMS system requires not only a removal from your main system but also a removal from your startup, login, and shutdown processes.

Perform the following steps to remove the CLI from your OpenVMS system:

1. Log on to your host system as a user with SYSLCK, SYSNAM, SYSPRV (or a system group UIC), TMPMBX, and CMKRNL privileges.
2. Type the following command at the command prompt to start the uninstallation process: `PRODUCT REMOVE IBMDSCLI`. A message similar to the following is displayed.

```
The following product has been selected:
  IBM AXPVMS DSCLI V5.0-1F96           Layered Product
Do you want to continue? [YES]
```

3. Press Enter to confirm the uninstallation. The following uninstallation confirmation message with completion status is displayed.

```
The following product will be removed from destination:
  IBM AXPVMS DSCLI V5.0-1F96           DISK$V732_ALPHA:[VMS$COMMON.]
Portion done: 0%...10%...20%...30%...40%...50%...60%...70%...80%...90%...100%
```

When the uninstallation process ends, a message similar to the following is displayed.

The following product has been removed:
IBM AXPVMS DSCLI V5.0-1F96

Layered Product

4. Remove the command-line interface startup, login, and shutdown functions from your system startup, login, and shutdown processes.

Removing the DS CLI from a System i system

This section contains information to help you uninstall the DS CLI from a System i system.

Because the DS CLI is installed on a System i system from a remote system, it is not possible to use the conventional DS CLI removal methods that you use with other systems.

When the DS CLI was installed onto your System i system, you used a remote system to perform the installation (for example, Windows, UNIX or AIX). Part of the installation process is the creation of an uninstaller. However, because you were using another system to do your installation, the uninstaller that was created was for the system that you installed from and not for the System i system. This uninstaller cannot be used to uninstall the DS CLI.

When you want to uninstall the DS CLI, you can use one of the following two methods:

- Uninstall directly from your i5/OS iSeries system.

You might use this method if you are not planning to upgrade the DS CLI and you want to totally remove the DS CLI from your System i system.

- Uninstall using a remote system.

You might use this method when you are upgrading the DS CLI, because after the removal, you can use this remote system to install the upgraded DS CLI.

Removing the DS CLI using your System i system directly

Complete this task to remove the DS CLI through the direct use of your System i system.

You cannot use the conventional DS CLI removal methods that are used on other systems because the installation of the DS CLI on your System i system was done from a remote system. The remote installation does not allow the creation of an uninstaller that can be used directly by your System i system for the removal process. However, it is possible to use your System i system directly (bypassing the uninstaller) to remove the DS CLI application.

You cannot use the uninstaller that was created for the DS CLI when you originally installed the DS CLI because it was created for the remote system that you used for the installation and not for the System i system.

You can perform this procedure at any time. However, it is common to perform this procedure when you want to remove the DS CLI from your system, but you do not intend to perform an associated upgrade of the DS CLI application.

Note: The i5/OS direct removal method requires that you use the i5/OS console mode and that you issue an i5/OS command. The following steps presume that you are logged in to the i5/OS and have the authority to issue a removal command.

Perform the following steps to remove the DS CLI through the direct use of your System i system:

1. Issue the following command from your i5/OS application:

```
RUNJAVA CLASS(run) PARM('-console')  
CLASSPATH('/QIBM/ProdData/Java400/jt400ntv.jar:/yourdir/_uninst  
/uninstall.jar')
```

Substitute your uninstall directory for *yourdir*.

2. Wait until the uninstall process is complete before you continue with your next process.

Removing the DS CLI from your System i system using the remote method

Complete this task to remove the DS CLI from your System i system using the remote method.

You cannot use the conventional DS CLI removal methods that are used on other systems because the installation of the DS CLI on your System i system has been done from a remote system. The remote installation does not allow the creation of an uninstaller that can be used directly by your System i system for the removal process. However, it is possible to use the remote removal method on your System i system to remove the DS CLI application.

Ensure that the remote system that you use to uninstall the DS CLI is network-attached to the System i system and is a supported platform for DS CLI.

You can use the following remote method to remove the DS CLI from a System i system. You can perform this procedure at any time. However, it is common to perform this procedure when you want to upgrade the DS CLI application, because the remote system that you are using to remove the DS CLI is typically the same system that you use for the upgrade.

To remove the DS CLI from your System i system using the remote method, perform the following steps:

1. Use FTP to copy the `uninstall.dat` and `uninstall.jar` files from the `uninstall` directory on your System i system into a common directory on the machine that you are using to do the uninstall.
2. Invoke the uninstall process from the common directory (or point to it using the `classpath [-cp]` option) by issuing the following command: `java -cp uninstall.jar run -os400 <system> <userid> <password>`
3. Wait until the uninstall process is completed before you continue to your next process.

Chapter 11. Troubleshooting

The topics in this chapter provide troubleshooting information related to your DS6000 series. Topics covered include analyzing, verifying, understanding, managing, and handling various problems.

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. Open a command prompt and navigate to the C:\Program Files\IBM\dsniserver\bin\ directory where the recovery tool (script) has been installed.
 2. Type the script name, `securityRecoveryUtility.bat -r`
 3. Press the **Enter** key. The script runs and the administrative account is unlocked. The password is reset to the default (admin).

Analyzing generic alert traps

The storage unit generates generic alert traps. The traps are sent by an SNMP agent to report asynchronous events.

Generic trap alert identifiers

coldStart (Generic 0)

Issued whenever the SNMP agent is reinitializing. Configuration data has changed.

warmStart (Generic 1)

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

authenticationFailure (Generic 4)

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

Analyzing service information messages for S/390 and zSeries systems

Service information messages (SIMs) are generated by a storage unit for S/390 and zSeries hosts. Before installation, you can use the customization work sheets to record the following service information: the SIM severity level and 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 DS6000 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 an 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

Logically removing a physical resource after physical removal

Follow these procedures to remove a physical resource from the DS Storage Manager after it has been removed from the server or expansion enclosure.

Problem

Follow these procedures if you have removed a DDM or expansion enclosure that was in use but have received an error message or log entry stating that the DDMs or expansion enclosures can no longer be recognized by the DS Storage Manager.

You are not required to perform these procedures, and you will not receive an error message, if you physically remove a DDM or expansion enclosure that is not in use. The resource will no longer appear in the list of physical resources in the summary tables.

Investigation

Follow these steps to physically replace the DDM and then properly take the DDM offline before physically removing it again:

1. Insert the DDM back into the enclosure.
2. Take the DDM offline through the Status page of the DS Storage Manager.
3. Once the DDM is offline, you can physically remove the DDM.

Follow these steps to remove the logical resources that were contained on the expansion enclosure:

1. Use the Modify Storage Unit -- Specify DDM packs page in the DS Storage Manager to logically remove all DDMs that are associated with the expansion enclosure.

2. Close the associated error log entry on the Logs page of the DS Storage Manager.

If, after you complete all of the steps above, you still receive an error message, call IBM Support.

IBM Support personnel can determine the cause of the problem and find the appropriate resolution.

Resolving IP address connectivity issues

Follow these procedures to regain connectivity to the server enclosure.

Problem

Follow these procedures to resolve IP address connectivity issues such as communications between the management console and the server enclosure. You might receive a communication error message or a timeout error log entry in the DS Storage Manager when you attempt to perform actions on the server enclosure.

Investigation

Verify lack of connectivity to the server enclosure

Use the **ping** command inside a command prompt or the Attempt connection page in the DS Storage Manager to attempt to retrieve IP address and network configuration information from the server enclosure. If you are able to retrieve the IP address and network information, attempt to perform the action again. If you are unable to retrieve the IP address and network configuration information, continue with the rest of these procedures.

Check the cabling between the management console and the server enclosure

Check that the cabling between the management console and the server enclosure has not become loose or dislodged.

Check that all of the cables are operational

If a cable develops a fault and you had connectivity prior to the error, you might receive an error entry or message stating that the connection to the server enclosure has been lost. Replace the specified cable to regain connectivity to the expansion enclosure.

If the DS Storage Manager still cannot locate the server enclosure

Follow the steps below to reset the IP address on the processor card in the server enclosure to the default IP address. Attempt to reconnect to the server enclosure using the default IP address after you complete these steps.

1. Use the serial conversion cable that was shipped with the enclosure to attach your computer to the serial port on the enclosure.
2. Use a terminal emulator to connect to the server enclosure through the serial port that is located on the processor card (for example, NetTerm or Windows HyperTerminal). If you are using Windows HyperTerminal, you might need to provide a connection name and icon before you can specify the connection settings. Choose a communications port, such as COM1 or COM3, and connect to the processor card using following settings:

Remote connection setting	Remote connection value
Bits per second	38400

Remote connection setting	Remote connection value
Data bits	8
Parity	None
Stop bits	1
Flow control	Hardware

3. Use the default user ID of guest to access the processor card. The ncnconf script begins automatically.

Note: Use the password that you assigned to the guest user ID when you set the initial IP addresses on the processor cards.

4. After the program has started, select Configure network parameters from the ncnconf Main Menu options.
5. To reset the IP address, perform the following steps:
 - a. Choose Reset to default static IP network configuration from the Network configuration menu options. You will receive a confirmation that the default IP address has been changed to the following defaults:
 - 172.30.143.213 for cluster0
 - 172.30.143.214 for cluster1

Note: The ncnconf program prevents you from setting an IP address and network mask combination that conflict with any of the following IP addresses:

- 172.30.143.213
 - 172.30.143.214
 - 192.0.2.0
 - 192.0.2.1
- b. Select Back to Main Menu to return to the ncnconf Main Menu.
 - c. Select Apply changes and exit from the options in the main menu to save your changes and exit the application.
6. After you are able to connect to the server enclosure, set the IP addresses on each processor card to meet your network needs.

If, after you complete all of the steps above, the DS Storage Manager cannot communicate with the server enclosure, call IBM Support.

IBM Support personnel can determine the cause of the problem and find a resolution that will regain connectivity to the storage.

Resolving power-on issues

Follow these procedures if you have problems powering on or off your server or expansion enclosures.

Problem

Follow these procedures if you unsuccessfully attempt to power on or off the server or expansion enclosure.

Investigation

If the enclosure is currently powered off and will not power on:

- Check to ensure that the storage unit is receiving power from the outlet.

Note: You can use the ac LED indicator on the power supply to determine if the power supply is receiving power from the outlet.
- Check the cord to the power supply.
Ensure that the cords are securely inserted into both the source outlet and the power supply.
- Ensure that the cords are not faulty.
Replace the power cord with a cord that you know is in working condition.
- If the problem is with an expansion enclosure, push the power button on the rear operator panel of the expansion enclosure.

If the enclosure is currently powered on and will not power off:

Ensure that all resources contained within the enclosure are offline. Use the Status page of the DS Storage Manager to verify that all resources that are contained within the enclosure are offline. Take online resources offline before reattempting to power off the enclosure.

If, after you complete all of the steps above, you still receive an error message, call IBM Support.

IBM Support personnel must determine the cause of the problem and find the appropriate resolution.

Verifying interenclosure connectivity

Follow these procedures to resolve interenclosure connectivity issues.

Problem

Follow these procedures if you receive error messages or log entries that state that the DS Storage Manager cannot recognize drives, arrays, or volumes that are located in an attached server or expansion enclosure. You can also use these procedures if you receive an error message or log entry that the DS Storage Manager cannot recognize a connected expansion enclosure. These procedures apply to both new volumes or enclosures that have been added as well as existing volumes or enclosures that are no longer recognized by the DS Storage Manager.

Investigation

Check interenclosure cabling

Check the cabling between the server enclosure and the expansion enclosure that contains the unrecognized storage. Verify that the cabling between the enclosures matches the cabling that is provided in the appropriate cabling diagram.

Check that all of the cables are operational

If a cable develops a fault, you will receive an error entry or message stating that the connection to that expansion enclosure has been lost. Replace the specified cable, small-form factor pluggable (SFP), or both, to regain connectivity to the expansion enclosure.

If, after you complete all of the steps above, storage resources are still not recognized by the DS Storage Manager, call IBM Support.

IBM Support personnel can determine the cause of the problem and find a resolution that will regain connectivity to the storage.

Verifying correct operation of the panels

Follow these procedures to determine if the front display panel or rear operator panel must be replaced.

Problem

Follow these steps if the server or expansion enclosure is powered on and operational, but the power-on indicator on either the front display panel or rear operator panel is not lit.

You can also follow these procedures if an error log entry appears in the Logs page of the DS Storage Manager and the system alert indicator on the front display panel or rear operator panel is not lit.

Investigation

If the enclosure is powered on but the power-on indicator is not lit

If the server or expansion enclosure is powered on, but the power-on indicator is not lit on either the front display panel or the rear operator panel, replace the faulty panel.

If an error log entry appears but the system alert indicator is not lit

If Logs page of the DS Storage Manager contains an error entry for a resource that is located within the server or expansion enclosure but the system alert indicator is not lit on either the front display panel or the rear operator panel, replace the faulty panel.

Use the Identify function to test the location LED indicators

Use the **Identify** button on the rear operator panel to locate connected expansion enclosures and to test the location LED indicators on the front display panel and rear operator panel.

If the location indicator on either the front display panel or rear operator panel does not light for a server enclosure, replace the faulty panel. If the location indicator on either the front display panel or rear operator panel does not light for an expansion enclosure that is correctly attached to the server enclosure, replace the faulty panel.

If, after you complete all of the steps above, you still receive an error message, call IBM Support.

IBM Support personnel can determine the cause of the problem and find the appropriate resolution.

Verifying host connectivity

Follow these procedures to resolve host connectivity issues.

Problem

These procedures apply to a situation where a host cannot connect to a storage unit.

Investigation

Check that the host is correctly configured in the DS Storage Manager

Use the Host systems pages in the DS Storage Manager to ensure that your host systems are configured correctly.

Check that the storage unit and hosts are both powered on.

The server enclosure, the connected expansion enclosures, and the attached host must be powered on before you can send information between the host and the storage unit.

Check that the cabling is correct between the hosts and the storage unit.

Miscabling between the storage unit and the host system can cause communication problems.

Check that the fiber optic cables and SFPs are operational.

Sequentially replace individual cables and small-form factor pluggables (SFPs) to determine if there are any cable problems.

If, after you complete all of the steps above, the storage unit still cannot communicate with the host system, call IBM Support.

IBM Support personnel can determine the cause of the problem and find a resolution that will reenable host communication.

Determining problems

This section contains information for determining problems with your storage unit hardware.

You can attempt the following sets of steps to determine hardware problems without consulting the Storage Manager:

- Performing a light path analysis
- Determining cable problems

Performing light path analysis

Complete this task to use the LED indicators on your storage unit to determine if there is a resource event that can be repaired without using the DS Storage Manager.

Use the LED indicators located on the front display panel and rear display panel as well as the individual resources within your server or expansion enclosure to determine if an event has occurred. Light path indicators provide indications of both fault and informational events.

1. Move to the front of your server enclosure.
2. Look at the front display panel to determine if any of the indicators are lit.
3. If the amber information indicator is lit, view the logs page of the DS Storage Manager to determine any further actions.
4. If the amber system alert light is lit, perform one of the following actions, depending on the other indicators that are lit:
 - If the fault on opposite side indicator is lit, move to the rear of the enclosure.
 - If the fault in external enclosure indicator is lit, move to the attached expansion enclosure that has the system alert indicator lit and repeat this step and the successive steps for that expansion enclosure.
 - If neither fault on opposite side nor the fault in external enclosure indicators are lit, stay at the front of the server enclosure. The fault is located on one of the disk drive modules.

5. Find the resource on the side of the enclosure that contain a fault that have a solid fault/service indicator. This is the resource that needs to be replaced.
6. Follow the removal and replacement procedures to replace the resource. After replacing the resource, the process automatically begins to bring the resource online. If, after the resource is fully online, the system alert indicators are still lit, repeat these steps to find the next resource that requires service.

Determining cable problems

Complete this task to determine if a cable is the cause of a resource fault event.

Before you can begin problem determination, you must have another cable of the same type as the cable that might be causing the fault event.

Note: For cabling information, see “Routing the cables” on page 53.

The following resources use cables to connect to other hosts, enclosures, networks, or power sources:

- Processor card
 - Power supply
1. Use the light path indicators or the DS Storage Manager to find the resource that has experienced a fault.
 2. If necessary, ensure that the external item that is connected to your enclosure, such as a host system or another enclosure, is offline.
 3. Disconnect the cable from the external item.
 4. Disconnect the cable from the server or expansion enclosure.
 5. Reconnect the replacement cable to the server or expansion enclosure.
 6. Reconnect the replacement cable to the external item.
 7. Follow the proper procedure to bring the external item online, if necessary.
 8. Watch the problem logs to see if the fault event occurs again.
 - If the fault event does not occur again, the cable was causing the problem.
 - If the fault event does occur again, the original cable was not the cause of the problem and you must replace the resource.
 9. Find and close the problem log entry for the resource through the DS Storage Manager.
 10. Bring the resource online.

Dispose of the cables that were determined to be the cause of the fault. Keep the cables that were not the cause of the fault for later use.

Sending problem determination data (real-time only)

Complete this task to collect a problem determination data file from the storage unit. You can send it to IBM technical support, save a copy to your local workstation, or delete it.

1. In the navigation, select **Real-time Manager, Manage Hardware**, and then **Storage units**.
2. Select the appropriate storage unit.
3. In the **Select Action** drop-down list, select **Copy and Send Problem Determination Data** and then **Go**. The Collect New PD Files page is displayed.

4. Select the type of information that you want to collect. You can select Traces, Dumps, or both file types. Select traces as the default option. The dumps option is primarily used if you are being instructed to do so by IBM technical support.
5. Type a description for the problem determination data file set that you are creating. This description is saved with the file set and can help distinguish the file set from other file sets. You must provide a description before you can proceed to the next step. You can edit this information in the Manage/Send Existing PD Files page.
6. Click **Copy**. The system collects and copies Traces (PE packages) and Dumps (statesaves) into files and displays them in the form of hyperlinks that can be managed from the Manage/Send Existing PD Files page. Click on the hyperlink to display additional information about the file set.

Note: You can also select file sets and click the **Save to Local disk** button or the **Delete** button.

7. Navigate to the Manage/Send Existing PD Files page. The table at the top of this page lists the problem determination data file sets that you created from the Collect New PD Files page.
8. Select the problem determination data file sets that you want to send to IBM technical support.
9. Click the **Send to IBM support** button to send the selected files to IBM technical support.

Note: To modify the description that you provided earlier, click on the file name hyperlink in the table and click **Edit description**.

Contacting IBM

Complete this task to contact IBM or view the IBM Support Web site.

You must have an Internet connection to contact IBM Support.

If you are contacting IBM Support for problem determination and resolution, you must also collect any information that can assist the support contact in diagnosing the issue. Such information can include, but is not limited to:

- Problem description
 - Machine serial numbers
 - Physical configuration information
 - Logical configuration information
 - Level of code that is installed on the system
 - Attached host types, host code levels, and world wide node names
 - Any applicable error messages
1. In the navigation, select **Real-time manager** → **Monitor System** → **Contact IBM**.
 2. Click the Contact IBM link on the page to open a new browser window for the IBM Support Web site.
 3. If you have not visited this site before, you must select your appropriate country from the drop down menu. If you have completed this step during a previous visit to the site, you do not have to select your country. The page opens automatically.

Providing PE package and state save information

This section contains information to help you use the DS CLI commands to provide your PE package and state save information to IBM. This information allows IBM to analyze your problem and to assist you directly with maintaining your DS6000 in an optimal state.

Note: The commands you issue for providing PE package and state save information can only be initiated from the DSCLI client installed locally on the management console connected to the target storage unit.

When you are directed by IBM Support personnel, you can use the DS CLI **mkpe**, **offloadss**, **sendpe**, and **sendss** commands to collect and to pass detailed information to IBM. You can use the following modes to provide this information to IBM:

- FTP
- An e-mail with the PE package, statesave file, or both, included as attachments
- CD sent by postal mail
- A manual method that does not use FTP, e-mail, or a CD sent by postal mail

Perform the following steps to send your pe package and statesave files to IBM support personnel:

1. Use the DS CLI **mkpe** and **offloadss** commands to initially collect and send pe package and state save information to IBM support personnel. See “Providing problem determination information using FTP” for additional information.
2. Use the DS CLI **sendpe** command to send the PE package to IBM Support after the initial FTP transfer failed. See “Sending the PE package after the FTP transfer has failed” on page 234 for additional information.
3. Use the DS CLI **sendss** command to send the state save information to IBM Support after the initial FTP transfer has failed. See “Sending the statesave file after the FTP transfer has failed” on page 236 for additional information.
4. Use the DS CLI **mkpe** command to collect and send an SMC only PE package to IBM support personnel. See “Providing a management console only PE package” on page 233
5. Use the alternative methods (e-mail or CD) to send PE package and state save information to IBM support personnel when there is no way you can use FTP. See “Providing problem determination information when FTP is not available” on page 237 for additional information.

Providing problem determination information using FTP

The problem determination information that you collect when you process the **mkpe** and **offloadss** commands is forwarded automatically using FTP to IBM support personnel. Because FTP is not a 100% reliable transport protocol, files that are sent by FTP are also saved on the storage management console. These saved files are automatically managed so that they do not accumulate and take up too much disk space on the storage management console.

To use the FTP feature of the **mkpe** and **offloadss** commands, ensure that you meet the following requirements:

- The **mkpe** and **offloadss** commands can only be initiated from the DSCLI client installed locally on the management console connected to the target storage unit.
- The server enclosures must be powered on and functioning.

- There must be an active IP connection between your storage management console and both processor cards.
- Adequate disk space must exist on your storage management console.
- The storage management console must be connected to the Internet, and its firewall must allow FTP to the IBM support Web address.
- You must have administrator authority to initiate and follow through with this process.

Note: The **mkpe** and **offloadss** commands only work when the DS CLI is installed on the storage management console PC. This is because the programs that gather and send the problem determination information to IBM, only exist on the storage management console PC.

The **mkpe** command collects the PE package data and presents the information in a summary format. The **offloadss** command collects the statesave data, which contains additional detailed information. Generally, when IBM support personnel request the PE package, they also request that you provide the statesave information. There are occasions, however, when only the PE package is needed for analysis.

Notes:

1. On a typical installation, MR1750_SM_HOME is set to C:\Program Files\IBM\DS6000StorageManager\SM.
2. Statesave files and PE packages are copied to the management console in subdirectory %MR1750_SM_HOME%\send2IBM before they are sent to IBM.
3. Statesave files and PE package files that have been sent by FTP to IBM can be found in the management console subdirectory, %MR1750_SM_HOME%\send2IBM\sent\.
4. The examples provided in this task are based on a Windows environment.

Perform the following steps to create PE package and statesave files and to *automatically* provide the requested information to IBM support personnel through the FTP process. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **mkpe** command to create the PE package file. Enter the **mkpe** command from the dscli prompt with the following parameters and variables:

```
dscli>mkpe -retry count storage_image_ID
```

Example

```
dscli>mkpe -retry 3 IBM.1750-75FA120
```

Note: Use the **-retry** parameter when you want the FTP process to attempt to send the files more than once. You can specify a value of 0 - 3, with 0 being the default value.

A successful process of this example command displays the following message:

```
The following PE package successfully generated and copied
from 1750-68FA120:
c:/Program_Files/IBM/DS6000StorageManager/SM/send2IBM/175051113AB15A.
IBM.0.NOPMH.20050328223544.c10.pe.zip
c:/Program_Files/IBM/DS6000StorageManager/SM/send2IBM/175051113AB15A.
IBM.0.NOPMH.20050328223807.c11.pe.zip
```

PE package successfully sent to IBM:
175051113AB15A.IBM.0.NOPMH.20050328223544.c10.pe.zip

PE package successfully sent to IBM:
175051113AB15A.IBM.0.NOPMH.20050328223807.c11.pe.zip

mkpe successfully completed.

Notes:

- a. The PE package information is collected on the management console subdirectory %MR1750_SM_HOME%\send2IBM and using FTP, is automatically sent to the IBM support Web address.
 - b. Files that are sent by FTP are automatically moved to the management console subdirectory, %MR1750_SM_HOME%\send2IBM\sent\.
2. Issue the **offloadss** command to create the statesave file. Enter the **offloadss** command at the dscli command prompt with the following parameters and variables:

```
dscli>offloadss -retry count storage_image_ID
```

Example

```
dscli>offloadss -retry 3 IBM.1750-68FA120
```

Note: Use the **-retry** parameter when you want the FTP process to attempt sending the files more than once. You can specify a value of 0 - 9, with 0 being the default value.

A successful process of this example command displays the following message:

```
offloadss: The following files offloaded from 1750-68FA120:  
/Program Files/IBM/DS6000StorageManager/SM/send2IBM/  
175051113AB15A.IBM.0.NOPMH.20050327243603.c10.dumptrace.tgz
```

```
offloadss: statesave FTPed to IBM:  
175051113AB15A.IBM.0.NOPMH.20050327243603.c10.dumptrace.tgz
```

```
offloadss: successfully completed.
```

Notes:

- a. The statesave information is collected on the management console subdirectory %MR1750_SM_HOME%\send2IBM and automatically sent to the IBM support Web address.
- b. The files that are sent by FTP are automatically moved to the management console subdirectory %MR1750_SM_HOME%\send2IBM\sent\.

If the automatic process that sends files by FTP is not available, you can use the **sendpe** or **sendss** commands at a later time. Or, contact IBM support personnel. You might be asked to use the following process:

1. Open a DOS window on the management console and change directories (cd) to the location of the files to be sent to IBM. For example, c:\program files\IBM\DS6000\SM\send2IBM
2. Type ftp testcase.software.ibm.com at the command prompt and press **Enter**. A prompt for *user* is displayed.
3. Type ftp at the user prompt and press **Enter**. The *password* prompt is displayed.

4. Type your e-mail address at the password prompt and press **Enter**. The command prompt is displayed.
5. Type `bin` and press **Enter**. It is important that you type *bin* to avoid the corruption of files.
6. Type `cd /ssd/toibm/sharkdumps` at the command prompt and press **Enter**. The command prompt at the `ssd/toibm/sharkdumps` directory is displayed.
7. Type `mput *.*` at the command prompt and press **Enter**. The file names that can be sent to IBM are displayed one at a time.
8. Answer *y* (yes) or *n* (no) to each file name that is presented to denote whether you want to transfer the file.

Notes:

1. You are not able to do a **dir** or **ls** command to validate that the file have been transferred.
2. You are not able to transfer the same file twice. If a situation arises where this needed, rename the file and then transfer it. Provide an e-mail to IBM support explaining this circumstance.

Providing a management console only PE package

Complete this task to create and send a management console (MC) only PE package to IBM support personnel. An MC only PE package contains the log files and configuration files that are associated only with the management console. It does not contain any files that are associated with the DS6000 nodes.

You must use the DS CLI **mkpe** command to collect and create the MC only PE package. To use this command and to use the automatic FTP feature to transfer the information to IBM support personnel, you must ensure that you meet the following requirements:

- The **mkpe** command can only be initiated from the DSCLI client installed locally on the management console connected to the target storage unit.
- The server enclosures must be turned on and functioning.
- There must be an active IP connection between your storage management console and both processor cards.
- Adequate disk space must exist on your storage management console.
- The storage management console must be connected to the Internet, and its firewall must allow FTP transfer to the IBM support Web address.
- You must have administrator authority to use the **mkpe** command.

Perform the following steps to create and send an MC only PE package file to IBM support personnel using the DS CLI **mkpe** command. The example commands that are displayed in this task are shown in two formats. The first format shows the type of information the command requires. The second format provides the command with declared values for the variables.

1. Call IBM support personnel and determine that a problem exists that affects only the management console. After you determine that you must send an MC only PE package to IBM, ensure that you also agree on a name to attach to the report before you hang up.

Note: The name is critical in identifying the PE package because there is no other way to determine which management console is being referenced by the PE package. If you do not have a name, you can assign one at the time that you initiate the command. Then, you must call the IBM

support personnel and inform them of the name that you have attached to the report before you can send it to them.

2. Issue the **mkpe** command to collect and create the MC only PE package. Enter the **mkpe** command at the dscli command prompt with the following parameters and variables:

```
dscli>mkpe -smonly -customer customer_name -retry count storage_image_ID
```

Example

```
dscli>mkpe -smonly -customer mine4results -retry 3 IBM.1750-75FA120
```

The output from this command example produces the following file name:

```
1750511SMC0000.mine4results.000.NOPMH.20051028223544.  
c10.SMCLog.zip
```

This file name is copied to the management console in subdirectory %MR1750_SM_HOME%\send2IBM. The FTP feature of this command automatically initiates and sends the MC only PE package to IBM support personnel. Based on the request that you made in the command (-retry 3), the FTP continues to initiate the transfer up to 3 times if it was not successful on the initial attempt.

When the FTP transfer is successful, the following two results occur:

- You receive a confirmation message that the transfer was successful like the following:
PE package successfully sent to IBM: 1750511SMC0000.mine4results.
000.NOPMH.20051028223544.c10.SMCLog.zip
- The MC only PE package file is moved into the management console subdirectory, %MR1750_SM_HOME%\send2IBM\sent\.

Note: If the FTP transfer was not successful, you can wait until later and attempt to resend the file using the **sendpe** command.

Sending the PE package after the FTP transfer has failed

Complete this task to send PE package files to IBM support personnel after the initial FTP transfer failed when you used the DS CLI **mkpe** command.

You must use the DS CLI **sendpe** command to send a PE package to IBM support personnel after the initial FTP transfer has failed when you used the DS CLI **mkpe** command. To use the **sendpe** command and to use the automatic FTP feature to transfer the information to IBM support personnel, you must ensure that you meet the following requirements:

- The **sendpe** command can only be initiated from the DSCLI client installed locally on the management console connected to the target storage unit.
- The server enclosures must be turned on and functioning.
- There must be an active IP connection between your storage management console and both processor cards.
- Adequate disk space must exist on your storage management console.
- The management console must be connected to the Internet, and its firewall must allow FTP transfer to the IBM support Web address.
- You must have administrator authority to use the **sendpe** command.

It is possible that when you initially created the PE package file and tried to send it using FTP that there was a problem with the connection. Or, it might be a case that your management console is located in a network that does not have direct access to the Internet or to IBM support personnel.

In any case, the time period for using the automatic feature of FTP that is associated with the DS CLI **mkpe** command has expired. However, you must still send the PE package to IBM support personnel. This can be done using the DS CLI **sendpe** command.

Perform the following steps to send a PE package file to IBM support personnel after the time period has expired for using the DS CLI **mkpe** command. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **lspe** command to locate the PE package files that are ready to be sent to IBM support personnel. Enter the **lspe** command at the dscli prompt with the following parameters and variables:

```
dscli>lspe -l -state ready
```

Example

```
dscli>lspe -l -state ready
```

Notes:

- a. **ready** is the actual value that you must state in the command.
 - b. Your search criteria is narrowed by using the **-state ready** parameter. Only the files that are in a **ready** state can be sent to IBM support personnel.
 - c. Use the **-l** parameter to provide PE package file details such as, file name, state (status), creation date, and size.
 - d. These files reside in the management console subdirectory `%MR1750_SM_HOME%\send2IBM` until they are sent to IBM support personnel.
2. Issue the **sendpe** command when you have determined which PE package file that you must send to the IBM support personnel. Enter the **sendpe** command at the dscli command prompt with the following parameters and variables:

```
dscli>sendpe -retry count pepackage_name
```

Example

```
dscli>sendpe -retry 3 175051113AB15A.IBM.0.NOPMH.20050328223544.  
c10.pe.zip
```

Notes:

- a. Use the **-retry count** parameter to specify the number of times that you want the command to try to initiate the FTP transfer process. You can specify a value of 0 - 9 attempts.
- b. If there is an FTP transfer problem that does not allow an FTP transfer, you can do one of the following:
 - Wait until a later time, use the **sendpe** command, and once again attempt to send the PE package.
 - Use an e-mail or copy the PE package file onto a CD and mail it to an address that is provided by IBM support personnel.

When the FTP transfer is successful, the following two results occur:

- You receive a confirmation message that the transfer was successful:

```
PE package successfully sent to IBM:  
175051113AB15A.IBM.0.NOPMH.20050328223544.  
c10.pe.zip
```

- The PE package file is moved into the management console subdirectory, `%MR1750_SM_HOME%\send2IBM\sent\`.

Sending the statesave file after the FTP transfer has failed

Complete this task to send statesave files to IBM support personnel after the initial FTP transfer has failed when you used the DS CLI **offloadss** command.

You must use the DS CLI **sendss** command to send a statesave file to IBM support personnel after the initial FTP transfer has failed when you used the DS CLI **offloadss** command. To use the **sendss** command and to use the automatic FTP feature to transfer the information to IBM support personnel, you must ensure that you meet the following requirements:

- The **sendss** command can only be initiated from the DSCLI client installed locally on the management console connected to the target storage unit.
- The server enclosures must be turned on and functioning.
- There must be an active IP connection between your storage management console and both processor cards.
- Adequate disk space must exist on your storage management console.
- The management console must be connected to the Internet, and its firewall must allow FTP transfer to the IBM support Web address.
- You must have administrator authority to use the **sendss** command.

It is possible that when you initially created the statesave file and tried to send it using FTP that there was a problem with the connection. Or, it might be a case that your management console is located in a network that does not have direct access to the Internet or to IBM support personnel.

In any case, the time period for using the automatic feature of FTP that is associated with the DS CLI **offloadss** command has expired. However, you must still send the pe package to IBM support personnel. This can be done using the DS CLI **sendss** command.

Perform the following steps to send a statesave file to IBM support personnel after the time period has expired for using the DS CLI **offloadss** command. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **lsss** command to locate the statesave files that are ready to be sent to IBM support personnel. Enter the **lsss** command at the dscli prompt with the following parameters and variables:

```
dscli>lsss -l -state ready
```

Example

```
dscli>lsss -l -state ready
```

Notes:

- a. **ready** is the actual value you must state in the command.
- b. Narrow your search criteria by using the **-state ready** parameter. These are the only files that can be sent to IBM support personnel.
- c. Use the **-l** parameter to provide the statesave file details such as, file name, state (status), creation date, and size.
- d. These files reside in the management console subdirectory `%MR1750_SM_HOME%\send2IBM` until they are sent to IBM support personnel.

2. Issue the **sendss** command when you have determined which PE package file that you must send to the IBM support personnel. Enter the **sendss** command at the dscli command prompt with the following parameters and variables:

```
dscli>sendss -retry count statesave_filename
```

Example

```
dscli>sendss -retry 3 175051113AB15A.IBM.0.NOPMH.050318180713.c10.dumptrace.tgz
```

Notes:

- a. Use the **-retry count** parameter to specify the number of times that you want the command to try to initiate the FTP transfer process. You can specify a value of 0 - 9 attempts.
- b. If there is an FTP transfer problem that does not allow an FTP transfer, you can do one of the following:
 - Wait until a later time, and then try again to send the statesave package with the **sendss** command.
 - Use an e-mail or copy the statesave file onto a CD and mail it to an address that is provided by IBM support personnel.

When the FTP transfer is successful, the following two results occur:

- You receive a confirmation message that the transfer was successful:

```
sends: Statesave sent to IBM:
175051113AB15A.IBM.0.NOPMH.050318180713.c10.dumptrace.tgz
sendss: successfully completed.
```
- The statesave file is moved into the management console subdirectory, `%MR1750_SM_HOME%\send2IBM\sent\`.

Providing problem determination information when FTP is not available

Complete this task to provide problem determination information to IBM if FTP is not available.

There might be times when you must provide problem determination information to IBM support personnel but your FTP connection is not available. You can still make use of the **mkpe** and **offloadss** commands to collect the information, but you must include the **-noftp** parameter. You can send the collected information to IBM support using e-mail or sending a CD that contains the information.

You must ensure that you are logged into the DS CLI application in interactive mode.

The **mkpe** command collects the PE package data and presents the information in a summary format. The **offloadss** command collects the statesave data, which contains additional detailed information. Generally, when IBM support personnel request the PE package, they also request that you provide the statesave information, as well. There are occasions however, when only the PE package is needed for analysis.

Perform the following steps to provide the requested information to IBM support:

Note: Obtain the address information from IBM Support where you can send your e-mail or CD.

1. Issue the **mkpe** command from the dscli prompt as follows: `dscli>mkpe -noftp IBM.1750-75FA120`. The PE package information is collected on the management console subdirectory `%MR1750_SM_HOME%\send2IBM`.

2. Create and send an e-mail with your attached PE package file, to the address IBM Support has directed that your file be sent. Or, send your CD by mail to the address that IBM Support has given you.

Note: The size of the attachment might exceed the limits that you have set for your e-mail attachments. It is suggested that the e-mail attachments be put into individual e-mails to preclude exceeding your limits. If there is still a problem, you might consider copying the files onto a CD and then mailing the CD to the address that IBM Support has given you.

3. Issue the **offloadss** command from the dscli prompt as follows: `dscli>offloadss -noftp IBM.1750-75FA120`. The statesave files are collected on the management console subdirectory `%MR1750_SM_HOME%\send2IBM`.
4. Create and send an e-mail with your attached statesave file to the address that IBM Support has directed that your file be sent. Or, send your CD by mail to the address that IBM Support has given you.
5. Wait for confirmation from IBM Support that they have received your e-mail and copy your statesave files and PE package files to a backup directory.

Note: When your problem has been resolved, you can delete the backup copy of your files.

6. Delete the PE package files from your backup directory.
7. Delete the statesave files from your backup directory.

Powering off

Complete this task to turn off the power to a server enclosure and all attached expansion enclosures.

Perform these steps to turn off the power to a storage unit:

1. Ensure that all I/O operations are stopped.
2. Quiesce all hardware resources in the server enclosure.
3. Press and hold the power button that is located on the left side of the rear operator panel on the server enclosure for five seconds. The unit completes the power-off sequence to turn off the server enclosure and each of the attached expansion enclosures. The server enclosure controls the power-off sequence for the entire storage unit, including all of the expansion enclosures. Each expansion enclosure turns off the enclosed resources in the necessary order. A delay of several minutes can occur before the system fully turns off all of the enclosures, depending on how many attached expansion enclosures are connected.

The enclosure is completely powered off after the fans have stopped spinning, the dc LED indicator on the power supply is flashing, and the ac LED indicator on the power supply is a solid green light. The power-on LED indicators on both the rear operator panel and the front display panel flash intermittently after the enclosure is fully turned off.

CAUTION:

Do not remove the ac power cords from any enclosure while any of the fans are still spinning. You must wait until the complete power off sequence has finished before you remove the power cords. Failure to follow the correct power off procedure may cause loss of access to your data or the inability to successfully restart the DS6000.

Chapter 12. Recovering

The topics in this chapter include concepts and procedures related to copy services and clusters.

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 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 DS6000 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 configuration, the process to create a consistent and restartable copy at your 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.

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 12 on page 241 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 12 provides an example of the implementation of failover and failback operations:

Table 12. 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 12. 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

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).

Returning production to Site A after planned and unplanned outages (failback)

Returning production to its original implementation is called a failback recovery. After restoring operations at Site A, 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. See Table 12 on page 241 for an example of the implementation of failover and failback operations.

Note: The process to move production back to your local site (Site A) for a planned outage is identical to the one that is used for an unplanned outage. Therefore, this procedure outlines the steps for both outages.

The failback operation resynchronizes the volumes in the following manner depending on the volume state:

- If a volume at Site A is in simplex state, all of the data for that volume is sent from Site B to Site A.
- If a volume at Site A is in full-duplex or suspended state and without changed tracks, only the modified data on the volume at Site B is sent to the volume at Site A.
- If a volume at Site A is in a suspended state but has tracks that have been modified, the volume at Site B will discover which tracks were modified at any site and send both 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. Site B volume is the source volume of the failback operation. This volume was initially the target volume of the relationship.

Note: 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 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). However, this process also works if the target volume is in simplex state.

Perform the following steps using the DS Storage Manager. You can also use the DS CLI to perform Copy Services functions.

1. Perform a failback recovery operation using volumes at Site B. This process copies all changed tracks from the target volumes back to the source volumes and copies over any tracks that were modified on the original source volumes.
2. Before returning to normal operation, quiesce your applications (still updating volumes at Site B) to cease all write I/O from updating the source volumes.

Note: On some host systems, such as AIX, Windows, and Linux, before performing FlashCopy operations, you must quiesce your applications that access FlashCopy source volumes. The source volumes must then be unmounted during the FlashCopy establishment. This is to ensure that there is no data in the buffers that could be flushed to the target volumes and potentially corrupt them. Depending on the host operating system, it might be necessary to unmount the source volumes.

3. From Site A, perform a failover recovery operation using the source volumes. This process converts the full-duplex target volumes at the Site A to suspended source volumes. The volumes at Site A start the change recording process while in failover mode.
4. Depending on your operating system, it might be necessary to rescan fibre-channel devices and mount the new source volumes at Site A.
5. From Site A, perform another failback recovery operation using the source volumes. This process resynchronizes the volumes at Site A with 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).
6. Mount your volumes at Site A and start your applications on your primary server.

Performing failover and failback operations (with Global Mirror)

Global Mirror is a long-distance, two-site, remote copy solution that uses asynchronous technology. Global Mirror is based on existing Global Copy and FlashCopy functions.

The following tasks are illustrated:

- Setting up your environment to use Global Mirror
- Using Global Mirror for a planned failover
- Using Global Mirror for an unplanned failover
- Using recovery procedures with and without the fast reverse feature

Global Mirror setup (DS Storage Manager)

Use this DS Storage Manager process to set up your environment to use Global Mirror. Global Mirror processing is used for general transactions as well as with data recovery transactions. When Global Mirror is used for data recovery it asynchronously copies data from a host to a recovery site and maintains data on a storage unit at the recovery site.

You can configure the following Global Mirror environment, which uses two sites (local and remote) and two or more storage units.

Local site

Contains A volumes (the source volume), which are copied to the recovery site using Global Copy

Recovery (or remote) site

Contains B volumes (the target volume and FlashCopy source volume) and C volumes (the FlashCopy target volume)

Note: A storage unit at the local site is designated as the Global Mirror master and all other local (or production) storage units are designated as subordinate storage units. The master storage unit sends commands to its subordinate storage units. These subordinates work together to create a consistency group and to communicate the FlashCopy commands to the recovery (or remote) site. All status is relayed back to the Global Mirror master.

You can perform this task using either the DS CLI or the DS Storage Manager.

DS Storage Manager setup steps

Perform the following steps to set up your environment for Global Mirror processing and to create a consistent backup copy of your data.

1. **Ensure that the storage units that you are using are configured, assigned, and operating in a normal state.** See "Storage Units — Main Page" for details.
2. **Ensure that fibre-channel paths are established.**
 - All Global Copy source and target pairs
 - Master and subordinate storage units

Note: To establish a volume pair, there must be a path between the LSSs in which the volumes reside. See "Creating remote mirror and copy paths" for the task procedure.

3. **Create Global Copy pairs.** Create the volumes from the local storage units to the recovery storage units using the path that you established. See "Creating Metro Mirror volume pairs" for the task procedure.

Note: When you create volume pairs using Global Copy, data on the volumes is not time consistent and generally not useful for the recovery site unless all of the data from the local site has been copied up to a specific point in time. To be useful for disaster recovery, the target volumes must periodically be made consistent and a copy of the data made. Be aware that the Global Copy source volumes are not active in the Global Mirror session until they have completed the first pass of their initial copy.

4. **Create FlashCopy relationships.** Create the relationships at the recovery site between the Global Copy secondary volumes and the FlashCopy target volumes. Ensure that you resynchronize all volumes that will be in the FlashCopy relationships. See "Creating a FlashCopy volume pair" for the task procedure. When creating the FlashCopy relationship, select the following options:

Enable Change Recording

Select this option to activate change recording on the volume pair participating in a FlashCopy relationship.

Note: The Persistent FlashCopy option is the default with this option.

Inhibit writes to target volume

Select this option to ensure that updates cannot be made to the target volume. This ensures data consistency on the target volume. If you select the Inhibit writes to target option, the change recording feature is not active on the target volume, only on the source volume.

Note: Do not select the **Initiate background copy** option so that data will only be copied from the source volume to the target volume if a track on the source volume is modified.

5. **Create your Global Mirror session.** See "Creating a new Global Mirror session" for the task procedure.
6. **Add volumes to your Global Mirror session.** See "Adding volumes to a Global Mirror session" for the task procedure.

After the setup is done, the following process describes how data consistency is created between the two sites:

1. Consistency groups of volumes are created at the local site.
2. Incremental copies of data are sent to the recovery site.
3. FlashCopy operation is performed at the recovery site.
4. Global Copy operations are resumed (only out-of-sync data is copied).
5. Steps are repeated, which are determined by the defined time periods.

Using Global Mirror for a planned failover and failback

Use this process to manage a planned failover involving two sites, which are referred to as local (or production) and remote (or recovery) sites, using Global Mirror.

During a planned outage, all storage units, servers, and networks are functional. Last and current states of all components are well-defined.

Examples for planned outages are system maintenance, disaster recovery tests, and training. The objective of the failover and failback procedures is to continue with consistent and current data (without any loss of transactions).

The procedure that is presented in this section makes the following assumptions:

- All storage units (local and remote) are functional and accessible.
- Applications are updating the primary volumes that are located at the local site.
- Paths are created from the local to the remote site.
- You have already created an environment to manage a planned outage.

Perform the following steps using the DS Storage Manager to begin the planned failover. You can also use the DS CLI.

Note: For details on individual tasks, see the related topics section.

1. Quiesce host I/O that is updating the primary volumes that is located in the local site.
2. Check the status of Global Mirror to ensure that Global Mirror is in a good state and that a consistency group has formed successfully at the desired interval.
3. Wait until two consistency groups form successfully.
4. Pause Global Mirror and suspend the Global Copy pairs (A->B).
5. Create paths from the B volumes to the A volumes.
6. Perform recovery failover on the B volumes (B->A).
7. Perform the fast reverse restore process from the C volumes to the B volumes, selecting the **Initiate background copy** option.

Notes:

- a. When you initiate the fast reverse restore process, Volume C becomes unusable.
 - b. There must be *no* I/O allowed to the B or C volumes during the fast reverse restore process.
8. Wait for the background copy to complete before continuing to the next step. The C to B FlashCopy relationship ends when the background copy completes.
 9. Initiate the FlashCopy from Volume B back to Volume C. Ensure that you also select the **Enable Change Recording** and **Inhibit writes to target volume** options. This creates a backup copy of the consistency group before applications begin to update the B volumes.
 10. Start the host I/O at the remote site on the B volumes. You will remain in production on the remote site in this configuration until you are ready to return production to the local site.
 11. When you are ready to return production to the local site, perform recovery failback (B->A) to resynchronize the A volumes. The application at the remote site remains active.
 12. After the resynchronize process has completed its first pass, quiesce the applications at your remote site so that the resynchronization can complete.

13. When the resynchronization has completed (no out-of-sync tracks), perform recovery failover and failback with Global Copy on Volume A to re-create the Volume A to Volume B Metro Mirror relationship.
14. Start the host I/O at the local site on the A volumes.
15. Resume the Global Mirror process.

Using Global Mirror for an unplanned failover and failback

Use this process to manage an unplanned failover and failback that involves two sites which are referred to as the local (or production) and remote (or recovery) sites.

Global Mirror provides two-site extended distance remote copy disaster recovery. When a disaster occurs at the local site, you must initiate the failover and failback recovery of consistent data on the remote site. Host activity can resume on the local site when the host recovers but not before a consistent set of data is copied to all primary volumes on the local site.

With Global Mirror, the data that the host writes to the storage unit at the local site is asynchronously shadowed to the storage unit at the remote site. A consistent copy of the data is then automatically maintained on the storage unit at the remote site.

Use of Global Mirror does not guarantee against data loss. During a disaster, data can only be restored to the last known consistent increment that was created. This means that data that is written to the primary site and is waiting to be transferred to the secondary site is lost whenever the two storage units can no longer communicate.

The following considerations apply when you use the Global Mirror recovery process:

- The Global Mirror master might still be running at the local site, especially if the disaster at the local site is a rolling disaster in which not all components fail simultaneously.
- The consistent copy at the remote site is *not* the secondary volume, but it is the FlashCopy target whose source is the secondary volume.
- Formation of a consistency group might have been in progress at the time of the failure.
- You can speed up recovery processing by choosing the **Fast Reverse** restore process that is explained later in this section.

Perform the following steps to use Global Mirror for an unplanned failover using the DS Storage Manager. You can also use the DS CLI for these steps.

Note: For details on individual tasks, see the related topics section.

1. Enter the Global Mirror session at the local site.

Note: Wait until the master storage unit completes the termination processing or enters the fatal state before continuing with the next step. Of course, this might not be possible if the local site has completely failed. If that is the case, proceed to the next step without waiting.

2. Issue a recovery failover request on the Global Copy volumes pair to force a stop of the Volume A to Volume B extended distance relationship and create a Volume B to Volume A Global Copy relationship.

Note: All B volumes must successfully process the recovery failover request before you can move to the next step.

3. Look at the session properties for Volume B and Volume C to ascertain the state of the consistency group between the B volume and C volume. You are looking primarily at the FlashCopy relationships and your analysis determines your next step in the recovery process. Act on your analysis as follows:
 - a. FlashCopy relationships are nonrestorable and all the sequence numbers are equal. No action to the consistency group is necessary.
 - b. FlashCopy relationships are restorable and all the sequence numbers are equal. Issue the FlashCopy **Discard changes** command to all the FlashCopy relationships in the consistency group.
 - c. All the FlashCopy sequence numbers are equal and at least one of the FlashCopy relationships is nonrestorable. Issue a FlashCopy **Commit changes** command to all the FlashCopy relationships in the consistency group that are restorable.
 - d. You have a mixed list of FlashCopy relationships, some are restorable and some are nonrestorable. The sequence numbers of the relationships that are restorable all have the same sequence number. The sequence numbers of the relationships that are nonrestorable are all equal, but they have a different number from the sequence number of those that are restorable. Issue a FlashCopy **Commit changes** command to all the FlashCopy relationships in the consistency group that are restorable.
 - e. You have a mixed list of FlashCopy relationships, some are restorable and some are nonrestorable. The sequence numbers are not the same within each type of relationship. The recovery plan cannot continue. The Global Mirror process has been corrupted. If the Global Mirror process has been corrupted you must recover your data using your last good backup.

Note: When the state of all the FlashCopy relationships are known, you might want to initiate a tape backup of Volume C.

4. Issue the fast reverse restore process from the C volumes to the B volumes, selecting the **Initiate background copy** option.

Note:

- a. When you initiate the fast reverse restore process, Volume C becomes unusable.
 - b. There must be *no* I/O allowed to the B or C volumes during the fast reverse restore process.
 - c. If you do not want to use the fast reverse restore process, use the Recovering from a disaster without using fast reverse restore procedure instead of this step.
5. Wait for the background copy to complete before continuing to the next step. The C to B FlashCopy relationship ends when the background copy completes.
 6. Initiate the FlashCopy from Volume B back to Volume C. Ensure that you also select the **Enable Change Recording** and **Inhibit writes to target volume** options. This creates a backup copy of the consistency group before applications begin to update the B volumes.
 7. Start the host I/O at the remote site on the B volumes. You will remain in production on the remote site in this configuration until you are ready to return production to the local site.

8. When you are ready to return production to the local site, perform the recovery failback (B->A) to resynchronize the A volumes. The application at the remote site remains active.
9. After the resynchronize process has completed its first pass, quiesce the applications at your remote site so that the resynchronization can complete.
10. When the resynchronization has completed (no out-of-sync tracks), perform recovery failover and failback with Global Copy on Volume A to re-create the Volume A to Volume B Metro Mirror relationship.
11. Start the host I/O at the local site on the A volumes.
12. Resume the Global Mirror process.

Recovering with and without the fast reverse restore feature

The fast reverse restore feature, when used with Global Mirror functions, allows a FlashCopy relationship to be reversed without waiting for the background copy of the original FlashCopy to complete.

You can reverse a FlashCopy relationship that has change recording enabled and restore the tracks from the target volume to the source volume. However, the background copy process of the original FlashCopy must complete before you can reverse the order of the FlashCopy relationship to its original source and target relationship.

With the fast reverse restore feature, you can reverse a FlashCopy relationship that has change recording enabled without waiting for the background copy of the original FlashCopy to complete. It is not required that you use the fast reverse restore feature during disaster recovery operations; however, the time to recover could be significantly longer if you do not.

Recovering from a disaster using the fast reverse process:

Use this process for the Global Mirror fast reverse restore feature, which allows a FlashCopy relationship to be reversed without waiting for the background copy of the original FlashCopy to complete.

You can use the fast reverse restore feature of FlashCopy to speed up the recovery process. Using the DS Storage Manager, you make the following selections:

1. Create a FlashCopy target on the existing Metro Mirror source by choosing the FlashCopy target and a source Metro Mirror .
2. Select the **Fast Reverse** restore option.

FlashCopy copies the partial data that had been copied to Volume C before the disaster occurred to Volume B. This creates a consistent copy on Volume B. Processing also provides a background copy of the consistent group for Volume B.

Note: When you initiate the fast reverse restore process, Volume C becomes unusable.

Recovering from a disaster without the fast reverse process:

Use this procedure in place of the fast reverse process. The fast reverse process allows a FlashCopy relationship to be reversed without waiting for the background copy of the original FlashCopy to complete.

It is not necessary to use the fast reverse feature of FlashCopy for your recovery. However, the wait for recovery is significantly longer if you do not.

1. Initiate a background copy on the Volume B to Volume C FlashCopy relationship.
2. Wait for the background copy to complete. Observe the out-of-sync tracks.
3. Select the **Reverse relationship** and **Initiate background copy** options on the Volume B to Volume C FlashCopy relationship. This reverses the FlashCopy relationship such that Volume C is the source and Volume B is the target.

Appendix. Using hexadecimal values

Use this conversion table to determine the hexadecimal value for a field or DS CLI command parameter that requires a hexadecimal value, or to determine the decimal value for a field or DS CLI command parameter that displays a hexadecimal value, such as the LSS and LUN fields or values.

Purpose

Table 13 provides a list of common decimal to hexadecimal conversion values. For values not listed in the table, you can use the Windows calculator to convert values. From the Windows calculator, click **View**, then **Scientific**. If you are converting from decimal to hexadecimal, ensure that **DEC** is selected in the row of radio buttons, enter a decimal number, and then select **HEX**. If you are converting from hexadecimal to decimal, ensure that **HEX** is selected in the row of radio buttons, enter a hexadecimal number, and then select **DEC**.

Table 13. Decimal to hexadecimal conversion

Decimal	Hexadecimal
1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9
10	A
11	B
12	C
13	D
14	E
15	F
16	10
17	11
18	12
19	13
20	14
21	15
22	16
23	17
24	18
25	19

Table 13. Decimal to hexadecimal conversion (continued)

26	1A
27	1B
28	1C
29	1D
30	1E
31	1F
32	20
64	40
100	64
128	80
200	C8
256	100
500	1F4
1000	3E8
8192	2000
10000	2710
64384	FB80
65280	FF00

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 DS6000 information:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. IBM Home Page Reader version 3.0 has been tested.
- You can operate features using the keyboard instead of the mouse.

Navigating by keyboard

You can use keys or key combinations to perform operations and initiate menu actions that can also be done through mouse actions. You can navigate the IBM System Storage DS6000 information from the keyboard by using the shortcut keys for your browser or Home Page Reader. See your browser Help for a list of shortcut keys that it supports. See the following Web site for a list of shortcut keys supported by Home Page Reader: http://www-306.ibm.com/able/solution_offerings/keyshort.html

Accessing the publications

You can find HTML versions of the IBM System Storage DS6000 information at the following Web site: <http://www.ehone.ibm.com/public/applications/publications/cgibin/pbi.cgi>

You can access the information using IBM Home Page Reader 3.0.

Notices

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been

estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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

Terms and conditions

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

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

- AIX
- DB2
- DFSMS/MVS
- DFSMS/VM
- e (logo)
- Enterprise Storage Server
- ES/9000
- ESCON
- FICON
- FlashCopy
- Graphically Dispersed Parallel Sysplex
- HACMP
- i5/OS
- IBM
- IntelliStation
- MVS/ESA
- Netfinity
- NetVista
- Operating System/400
- OS/400
- RS/6000
- S/390
- Seascape
- SNAP/SHOT
- SP
- System/390
- System p5
- System Storage
- Versatile Storage Server
- Virtualization Engine
- VSE/ESA
- z/Architecture
- z/OS
- z/VM
- zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

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

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

Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.

Federal Communications Commission (FCC) statement

This equipment has been tested and complies with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the users authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Industry Canada compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conform à la norme NMB-003 du Canada.

European community compliance statement

This product is in conformity with the protection requirements of EC Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

Germany only

Zulassungsbescheinigung laut Gesetz ueber die elektromagnetische Vertraeglichkeit von Geraeten (EMVG) vom 30. August 1995.

Dieses Geraet ist berechtigt, in Uebereinstimmung mit dem deutschen EMVG das EG-Konformitaetszeichen - CE - zu fuehren.

Der Aussteller der Konformitaetserklaeung ist die IBM Deutschland.

Informationen in Hinsicht EMVG Paragraph 3 Abs. (2) 2:

Das Geraet erfuellt die Schutzanforderungen nach EN 50082-1 und EN 55022 Klasse A.

EN 55022 Klasse A Geraete beduerfen folgender Hinweise:

Nach dem EMVG:

"Geraete duerfen an Orten, fuer die sie nicht ausreichend entstoert sind, nur mit besonderer Genehmigung des Bundesministeriums fuer Post und Telekommunikation oder des Bundesamtes fuer Post und Telekommunikation betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Stoerungen zu erwarten sind." (Auszug aus dem EMVG, Paragraph 3, Abs.4)

Dieses Genehmigungsverfahren ist nach Paragraph 9 EMVG in Verbindung mit der entsprechenden Kostenverordnung (Amtsblatt 14/93) kostenpflichtig.

Nach der EN 55022:

"Dies ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funkstoerungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Massnahmen durchzufuehren und dafuer aufzukommen."

Anmerkung:

Um die Einhaltung des EMVG sicherzustellen, sind die Geraete wie in den Handbuechern angegeben zu installieren und zu betreiben.

Japanese Voluntary Control Council for Interference (VCCI) class A statement

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Korean Ministry of Information and Communication (MIC) statement

Please note that this device has been certified for business use with regard to electromagnetic interference. If you find this is not suitable for your use, you may exchange it for one of residential use.

Taiwan class A compliance statement

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

VS07171L

Glossary

This glossary includes terms for the IBM 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://www-306.ibm.com/ibm/terminology/index.html>. Definitions derived from this source have the symbol (GC) after the definition.
- *The Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions derived from this book have the symbol (I) after the definition. Definitions taken from draft international standards, committee drafts, and working papers that the ISO/IEC JTC1/SC1 is developing have the symbol (T) after the definition, indicating that final agreement has not been reached among the participating National Bodies of SC1.

This glossary uses the following cross-reference forms:

- See** Refers the reader to one of two kinds of related information:
- A term that is the expanded form of an abbreviation or acronym. This expanded form of the term contains the full definition.
 - A synonym or more preferred term

See also Refers the reader to one or more related terms.

Contrast with

Refers the reader to a term that has an opposite or substantively different meaning.

Numerics

- 750** A model of the Enterprise Storage Server featuring a 2-way processor with limited physical storage capacity. This model can be updated to the model 800.
- 800** A model of the Enterprise Storage Server featuring a standard processor or an optional Turbo processor. The Model 800 supports RAID 5, RAID 10, and 15000 rpm drives. Model 800 supersedes Model F20.
- 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** The machine type for the IBM System Storage DS8000 series. Models for the DS8000 series include base units and expansion units .
- 3390** The machine number of an IBM disk storage system. The Enterprise Storage Server, when interfaced to IBM zSeries hosts, is set up to appear as one or more 3390 devices, with a choice of 3390-2, 3390-3, or 3390-9 track formats.
- 3990** The machine number of an IBM control unit.
- 7133** The machine number of an IBM disk storage system. The Model D40 and 020 drawers of the 7133 can be installed in the 2105-100 expansion enclosure of the ESS.

A

access 1) To obtain 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 zSeries, a relationship that is created between a device and one or more channel paths during the processing of certain conditions. See also *implicit allegiance*, *contingent allegiance*, and *reserved allegiance*.

allocated storage

The space that is allocated to volumes but not yet assigned. Contrast with *assigned storage*.

American National Standards Institute (ANSI)

An organization of producers, consumers, and general interest groups that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States. (A)

anonymous

In the DS Storage Manager, the label on an icon that represents all connections that are using fibre-channel adapters between the storage unit and hosts but are not completely defined to the storage unit. See also *anonymous host*, *pseudo host*, and *access-any mode*.

anonymous host

Synonym for *pseudo host*. Contrast with *anonymous* and *pseudo host*.

ANSI See *American National Standards Institute*.

APAR See *authorized program analysis report*. (GC)

API See *application programming interface*.

application programming interface

An interface that allows an application program that is written in a high-level language to use specific data or functions of the operating system or another program.

arbitrated loop

A fibre-channel topology that enables the interconnection of a set of nodes. See also *point-to-point connection* and *switched fabric*.

array An ordered collection, or group, of physical devices (disk drive modules) that is used to define logical volumes or devices. In the storage unit, an array is a group of disks that the user designates to be managed by the RAID technique. See also *redundant array of independent disks*.

ASCII (American National Standard Code for Information Interchange) The standard code, using a coded character set consisting of 7-bit coded characters (8 bits including parity check), that is used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. (A) Some organizations, including IBM, have used the parity bit to expand the basic code set.

assigned storage

The space that is allocated to a volume and that is assigned to a port.

authorized program analysis report (APAR)

A request for correction of a defect in a current release of an IBM-supplied program. (GC)

availability

The degree to which a system or resource is capable of performing its normal function. See *data availability*.

B

bay The physical space that is used for

installing SCSI and fibre-channel host adapter cards. The DS6000 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 DS6000 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

For zSeries, the logical functions that a system requires to perform the function of a channel subsystem.

CKD See *count key data*.

CLI See *command-line interface*. See also *Copy Services command-line interface*.

cluster

1) A partition capable of performing all DS6000 functions. With two clusters in the DS6000, 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 DS6000. 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 *Copy Services command-line interface*.

compact disc

An optically read disc, typically storing approximately 660 MB. CD-ROM (compact disc read-only memory) refers to the read-only format used to distribute DS6000 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 DS6000 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 zSeries, a storage server with ESCON or OEMI interfaces. The control unit adapts a native device interface to an I/O interface that a zSeries host system supports.
- 3) The portion of the storage unit that supports the attachment of emulated count key data devices over FICON or OEMI interfaces. See also *cluster*.

control-unit-initiated reconfiguration (CUIR)

A software mechanism that the DS6000 uses to request that an operating system of a zSeries host verify that one or more subsystem resources can be taken offline for service. The DS6000 can use this process to automatically vary channel paths offline and online to facilitate bay service or concurrent code installation. Depending on the operating system, support for this process might be model dependent, might depend on the IBM TotalStorage Enterprise Storage Server Subsystem Device Driver, or might not exist.

Coordinated Universal Time (UTC)

The international standard of time that is kept by atomic clocks around the world.

Copy Services

A collection of optional software features, with a Web-browser interface, used for configuring, managing, and monitoring data-copy functions.

Copy Services CLI

See *Copy Services command-line interface*.

Copy Services domain

See *Copy Services server group*.

Copy Services client

Software that runs on each DS6000 cluster in the Copy Services server group and that performs the following functions:

- Communicates configuration, status and connectivity information to the Copy Services server
- Performs data-copy functions on behalf of the Copy Services server

Copy Services command-line interface (Copy Services CLI)

The command-line interface software that is provided with DS6000 Copy Services and used for invoking Copy Services functions from host systems attached to the DS6000. See also *command-line interface*.

Copy Services server

A cluster that the Copy Services administrator designates to perform the DS6000 Copy Services functions.

Copy Services server group

A collection of user-designated DS6000 clusters participating in Copy Services functions that a designated, active, Copy Services server manages. A Copy Services server group is also called a Copy Services domain. See also *domain*.

count field

The first field of a count key data (CKD) record. This eight-byte field contains a four-byte track address (CCHH). It defines the cylinder and head that are associated with the track, and a one-byte record number (R) that identifies the record on the track. It defines a one-byte key length that specifies the length of the record's key field (0 means no key field). It defines a two-byte data length that specifies the length of the record's data field (0 means no data field). Only the end-of-file record has a data length of zero.

count key data (CKD)

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 zSeries volume types: 3390-2, 3390-3, 3390-9, 3390-2 (3380-track mode), or 3390-3 (3380-track mode). See also *count-key-data*, *interleave*, *standard volume*, and *volume*.

CUT See *Coordinated Universal Time*.

cyclic redundancy check (CRC)

A redundancy check in which the check key is generated by a cyclic algorithm. (T)

cylinder

A unit of storage on a CKD device with a fixed number of tracks.

D

DA See *device adapter*.

daisy chain

See *serial connection*.

DASD

See *direct access storage device*.

DASD fast write (DFW)

A function of a storage server in which active write data is stored in nonvolatile cache, thus avoiding exposure to data loss.

data availability

The degree to which data is available when needed, typically measured as a percentage of time that the system would be capable of responding to any data request (for example, 99.999% available).

data compression

A technique or algorithm used to encode data such that the encoded result can be stored in less space than the original data. The original data can be recovered from the encoded result through a reverse technique or reverse algorithm. See also *compression*.

Data Facility Storage Management Subsystem (DFSMS)

An operating environment that helps automate and centralize the management of storage. To manage storage, DFSMS provides the storage administrator with control over data class, storage class, management class, storage group, and automatic class selection routine definitions.

data field

The optional third field of a count key data (CKD) record. The count field specifies the length of the data field. The data field contains data that the program writes.

data record

The basic unit of zSeries storage on a DS6000, also known as a count-key-data (CKD) record. Data records are stored on a track. The records are sequentially numbered starting with 0. The first record, R0, is typically called the track descriptor record and contains data that the operating system normally uses to manage the track. See also *count-key-data* and *fixed-block architecture*.

data set FlashCopy

For zSeries hosts, a feature of FlashCopy that indicates how many partial volume FlashCopy relationships are active on a volume.

data sharing

The ability of multiple host systems to concurrently utilize data that they store on one or more storage devices. The storage unit enables configured storage to be accessible to any, or all, attached host systems. To use this capability, the host

program must be designed to support data that it is sharing.

DDM See *disk drive module*.

DDM group
See *disk pack*.

dedicated storage
Storage within a storage unit that is configured such that a single host system has exclusive access to the storage.

demote
To remove a logical data unit from cache memory. A storage server demotes a data unit to make room for other logical data units in the cache or because the logical data unit is not valid. The storage unit must destage logical data units with active write units before they can be demoted. See also *destage*.

destage
To move data from an online or higher priority to an offline or lower priority device. The storage unit stages incoming data into cache and then destages it to disk.

device For zSeries, a disk drive.

device adapter (DA)
A physical component of the DS6000 that provides communication between the clusters and the storage devices. The DS6000 has eight device adapters that it deploys in pairs, one from each cluster. Device adapter pairing enables the DS6000 to access any disk drive from either of two paths, providing fault tolerance and enhanced availability.

device ID
The unique two-digit hexadecimal number that identifies the logical device.

device interface card
A physical subunit of a storage cluster that provides the communication with the attached device drive modules.

device number
For zSeries, a four-hexadecimal-character identifier, for example 13A0, that the systems administrator associates with a device to facilitate communication between the program and the host operator. The device number is associated with a subchannel.

device sparing
A subsystem function that automatically copies data from a failing device drive module to a spare device drive module. The subsystem maintains data access during the process.

DFS See *distributed file service*.

DFSMS
See *Data Facility Storage Management Subsystem*.

direct access storage device (DASD)
1) A mass storage medium on which a computer stores data.
2) A disk device.

disk cage
A container for disk drives. Each disk cage supports eight disk packs (64 disks).

disk drive
Standard term for a disk-based nonvolatile storage medium. The DS6000 uses hard disk drives as the primary nonvolatile storage media to store host data.

disk drive module (DDM)
A field replaceable unit that consists of a single disk drive and its associated packaging.

disk drive module group
See *disk pack*.

disk 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.
- 3) A Copy Services server group, in other words, the set of clusters the user designates to be managed by a particular Copy Services server.

domain name system (DNS)

In TCP/IP, the server program that supplies name-to-address translation by mapping domain names to internet addresses. The address of a DNS server is the internet address of the server that hosts the DNS software for the network.

dotted decimal notation

A convention used to identify IP addresses. The notation consists of four 8-bit numbers written in base 10. For example, 9.113.76.250 is an IP address that contains the octets 9, 113, 76, and 250.

drawer

A unit that contains multiple device drive modules and provides power, cooling, and related interconnection logic to make the device drive modules accessible to attached host systems.

- drive**
- 1) A peripheral device, especially one that has addressed storage media. See also *disk drive module*.
 - 2) The mechanism used to seek, read, and write information on a storage medium.

DS6000

See *IBM System Storage DS6000*.

DS6000 Batch Configuration tool

A program that automatically configures a DS6000. 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.

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*.

EPO See *emergency power off*.

ERDS See *error-recording data set*.

error-recording data set (ERDS)

On zSeries hosts, a data set that records data-storage and data-retrieval errors. A service information message (SIM) provides the error information for the ERDS.

error recovery procedure

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*.

ESD See *electrostatic discharge*.

eserver

See *IBM e(logo)server*.

ESSNet

See *IBM TotalStorage Enterprise Storage Server Network*.

express configuration

A method for configuring a storage complex, where the storage server simplifies the task by making several configuration decisions for the user.

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 DS6000 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 DS6000 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 zSeries command protocols). Refer to American National Standards Institute (ANSI) X3.230-199x. See also *Fibre Channel Protocol*.

fibre-channel topology

An interconnection topology supported on fibre-channel adapters. See also *point-to-point connection*, *switched fabric*, and *arbitrated loop*.

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 DS6000. 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 DS6000 that can make an instant copy of data, that is, a point-in-time copy of a volume.

FlashCopy relationship

A mapping of a FlashCopy source volume and a FlashCopy target volume that allows a point-in-time copy of the source volume to be copied to the target volume. FlashCopy relationships exist from the time that you initiate a FlashCopy operation until the storage unit copies all data from the source volume to the target volume or until you delete the FlashCopy relationship, if it is persistent.

FRU See *field replaceable unit*.

FTP See *File Transfer Protocol*.

full duplex
See *duplex*.

fuzzy copy
A function of the Global Copy feature wherein modifications to the primary logical volume are performed on the secondary logical volume at a later time. The original order of update is not strictly maintained. See also *Global Copy*.

G

GB See *gigabyte*.

GDPS See *Geographically Dispersed Parallel Sysplex*.

Geographically Dispersed Parallel Sysplex (GDPS)

A zSeries multisite application-availability solution.

gigabyte (GB)
A gigabyte of storage is 10⁹ bytes. A gigabyte of memory is 2³⁰ bytes.

Global Copy
An optional capability of the DS6000 remote mirror and copy feature that maintains a fuzzy copy of a logical volume on the same DS6000 or on another DS6000. 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 DS6000 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 an AS/400 or iSeries host that enables the user to display and work with system hardware resources and to debug input-output processors (IOP), input-output adapters (IOA), and devices.

HCD See *Hardware Configuration Data*.

HDA See *head 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 DS6000. 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, zSeries) or of the open-systems type, that is connected to the DS6000. Hosts are connected through 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*.

I

i5/OS The IBM operating system that runs the IBM i5/OS and e(logo)Server i5 server families of servers.

IBM e(logo)server

The IBM brand name for a series of server products that are optimized for e-commerce. The products include the iSeries, pSeries, xSeries, and zSeries.

IBM product engineering (PE)

The third-level of IBM service support. Product engineering is composed of IBM engineers who have experience in supporting a product or who are knowledgeable about the product.

IBM Serial Storage adapter

A physical adapter based on the IBM Serial Storage architecture. IBM Serial Storage adapters connect disk drive modules to DS6000 clusters.

IBM System Storage

The brand name used to identify storage products from IBM, including the IBM System Storage DS6000. See also *IBM System Storage DS6000* and *IBM System Storage DS Storage Manager*.

IBM System Storage DS6000

A member of the IBM System Storage Resiliency Family of storage servers and attached storage devices (disk drive

- modules). The DS6000 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 is specific to the DS6000.
- IBM System Storage DS Storage Manager (DS Storage Manager)**
Software with a Web-browser interface for configuring the DS6000.
- IBM TotalStorage Enterprise Storage Server Network (ESSNet)**
A private network providing Web browser access to the Enterprise Storage Server. IBM installs the ESSNet software on an IBM workstation called the IBM TotalStorage ESS Master Console, supplied with the first ESS delivery.
- IBM System Storage Management Console (MC)**
An IBM workstation that acts as the focal point for configuration, Copy Services management, and maintenance for the DS6000. It includes a Web browser that provides links to the user interface, including the DS Storage Manager and the DS6000 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 DS6000 and the IBM TotalStorage Enterprise Storage Server, Models 750 and 800.
- IML** See *initial microcode load*.
- implicit allegiance**
In Enterprise Systems Architecture/390, a relationship that a control unit creates between a device and a channel path when the device accepts a read or write operation. The control unit guarantees access to the channel program over the set of channel paths that it associates with the allegiance.
- initial microcode load (IML)**
The action of loading microcode for a computer into that computer's storage.
- initial program load (IPL)**
The action of loading software into a computer, typically an operating system that controls the computer.
- initiator**
A SCSI device that communicates with and controls one or more targets. Contrast with *target*.
- i-node** The internal structure in an AIX operating system that describes the individual files in the operating system. It contains the code, type, location, and owner of a file.
- input/output (I/O)**
Pertaining to (a) input, output, or both or (b) a device, process, or channel involved in data input, data output, or both.
- input/output configuration data set**
A configuration definition built by the I/O configuration program (IOCP) and stored on disk files associated with the processor controller.
- interleave**
To automatically create two striped partitions across the drives in a RAID-5 array, both of which use the count-key-data (CKD) record format.
- Internet Protocol (IP)**
In the Internet suite of protocols, a protocol without connections that routes data through a network or interconnecting networks and acts as an intermediary between the higher protocol layers and the physical network. The upper layer supports one or more logical protocols (for example, a SCSI-command protocol and a zSeries command protocol). Refer to ANSI X3.230-199x. The IP acronym is the IP in TCP/IP. See also *Transmission Control Protocol/Internet Protocol*.

invalidate

To remove a logical data unit from cache memory because it cannot support continued access to the logical data unit on the device. This removal might be the result of a failure within the storage server or a storage device that is associated with the device.

I/O See *input/output*.

I/O adapter (IOA)

An input-output adapter on the PCI bus.

IOCDs

See *input/output configuration data set*.

IOCP See *I/O Configuration Program*.

I/O Configuration Program (IOCP)

A program that defines to a system all the available I/O devices and channel paths.

I/O device

An addressable read and write unit, such as a disk drive device, magnetic tape device, or printer.

I/O interface

An interface that enables a host to perform read and write operations with its associated peripheral devices.

I/O Priority Queueing

A facility in the Workload Manager of zSeries that enables the system administrator to set priorities for queueing I/Os from different system images. See also *multiple allegiance* and *parallel access volume*.

I/O processor (IOP)

Controls input-output adapters and other devices.

I/O sequential response time

The time an I/O request is queued in processor memory waiting for previous I/Os to the same volume to complete.

IP See *Internet Protocol*.

IPL See *initial program load*.

iSeries

An IBM e(logo)server product that emphasizes integration. It is the successor to the AS/400 family of servers.

J**Java Virtual Machine (JVM)**

A software implementation of a central

processing unit (CPU) that runs compiled Java code (applets and applications). (GC)

JVM See *Java Virtual Machine*.

K

KB See *kilobyte*.

key field

The second (optional) field of a count key data record. The key length is specified in the count field. The key length determines the field length. The program writes the data in the key field and uses the key field to identify or locate a given record. The subsystem does not use the key field.

kilobyte (KB)

1) For processor storage, real, and virtual storage, and channel volume, 2^{10} or 1024 bytes.

2) For disk storage capacity and communications volume, 1000 bytes.

Korn shell

Interactive command interpreter and a command programming language.

KPOH

See *thousands of power-on hours*.

L

LAN See *local area network*.

last-in first-out (LIFO)

A queuing technique in which the next item to be retrieved is the item most recently placed in the queue. (A)

LBA See *logical block address*.

LCU See *logical control unit*.

least recently used (LRU)

1) The algorithm used to identify and make available the cache space that contains the least-recently used data.

2) A policy for a caching algorithm that chooses to remove from cache the item that has the longest elapsed time since its last access.

LED See *light-emitting diode*.

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.

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 DS6000 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 DS6000) associated with the processing of I/O operations directed to a single host-accessible emulated I/O device. The associated storage is referred to as a logical volume. The logical device is mapped to one or more host-addressable units, such as a device on a zSeries I/O interface or a logical unit on a SCSI I/O interface, such that the host initiating I/O operations to the I/O-addressable unit interacts with the storage on the associated logical device.

logical partition (LPAR)

For zSeries, a set of functions that create the programming environment in which more than one logical partition (LPAR) is established on a processor. An LPAR is conceptually similar to a virtual machine environment except that the LPAR is a function of the processor. Also, the LPAR

does not depend on an operating system to create the virtual machine environment. (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 DS6000 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 DS6000 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 DS6000 and might have other implications relative to the operation of certain functions. There is a one-to-one mapping between a CKD logical subsystem and a zSeries control-unit image.

For zSeries hosts, a logical subsystem represents a logical control unit (LCU). Each control-unit image is associated with only one logical subsystem. See also *control-unit image*.

logical unit

In open systems, a logical disk drive.

logical unit number (LUN)

In the SCSI protocol, a unique number that is used on a SCSI bus to enable it to differentiate between separate devices, each of which is a logical unit.

logical volume

The storage medium that is associated with a logical disk drive. A logical volume typically resides on one or more storage devices. The DS6000 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 DS6000 uses for locating errors. The LRC checks the data as it progresses from the host, through the DS6000 controller, into the device adapter, and to the array.

longwave laser adapter

A connector that is used between a host and the DS6000 to support longwave fibre-channel communication.

loop The physical connection between a pair of device adapters in the DS6000. 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 DS6000, the transfer of data from one storage unit to another, such as from a 3390 to the DS6000.

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 DS6000 Multipath Subsystem Device Driver*.

multiple allegiance

A DS6000 hardware function that is independent of software support. This function enables multiple system images to concurrently access the same logical volume on the DS6000 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 DS6000 that creates backup copies from one source to multiple targets by simultaneously establishing multiple FlashCopy relationships.

multiple virtual storage (MVS)

Implies MVS/390, MVS/XA, MVS/ESA, and the MVS element of the zSeries operating system.

multiplex

The action of transmitting simultaneously.

multiport serial adapter (MSA)

An adapter on the IBM System Storage Management Console that has multiple ports to which a DS6000 can be attached.

multiprocessor

A computer that includes two or more processors that have common access to a main storage. For the DS6000, the multiprocessors operate in parallel.

MVS See *multiple virtual storage*.

N

name server

A server that stores names of the participating DS6000 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 DS6000 is a node in a fibre-channel network.

non-RAID

A disk drive set up independently of other disk drives and not set up as part of a disk pack to store data using the redundant array of disks (RAID) data-striping methodology.

nonremovable medium

A recording medium that cannot be added to or removed from a storage device.

nonvolatile storage (NVS)

Memory that stores active write data to avoid data loss in the event of a power loss.

NVS See *nonvolatile storage*.

O

octet

In Internet Protocol addressing, one of the four parts of a 32-bit integer presented in dotted decimal notation. See also *dotted decimal notation*.

OEMI See *original equipment manufacturer's information*.

open system

A system whose characteristics comply with standards made available

throughout the industry and that therefore can be connected to other systems complying with the same standards. Applied to the DS6000, such systems are those hosts that connect to the DS6000 through SCSI or FCP protocols. See also *small computer system interface* and *Fibre Channel Protocol*.

organizationally unique identifier (OUI)

An IEEE-standards number that identifies an organization with a 24-bit globally unique assigned number referenced by various standards. OUI is used in the family of 802 LAN standards, such as Ethernet and Token Ring.

original equipment manufacturer's information (OEMI)

A reference to an IBM guideline for a computer peripheral interface. The interface uses ESA/390 logical protocols over an I/O interface that configures attached units in a multidrop bus topology.

OS/390

The IBM operating system that includes and integrates functions that many IBM software products (including the MVS operating system) previously provided for the IBM S/390 family of enterprise servers.

OUI See *organizationally unique identifier*.

P

panel The formatted display of information that appears on a display screen.

parallel access volume (PAV)

An advanced function of the DS6000 that enables OS/390 and z/OS systems to issue concurrent I/O requests against a count key data logical volume by associating multiple devices of a single control-unit image with a single logical device. Up to eight device addresses can be assigned to a PAV. The PAV function enables two or more concurrent write operations to the same logical volume, as long as the write operations are not to the same extents. See also *extent*, *I/O Priority Queueing*, and *multiple allegiance*.

parity A data checking scheme used in a computer system to ensure the integrity

of the data. The RAID implementation uses parity to re-create data if a disk drive fails.

path group

In zSeries architecture, a set of channel paths that are defined to a control unit as being associated with a single logical partition (LPAR). The channel paths are in a group state and are online to the host. See also *logical partition*.

path group identifier

In zSeries architecture, the identifier that uniquely identifies a given logical partition (LPAR). The path group identifier is used in communication between the LPAR program and a device. The identifier associates the path group with one or more channel paths, thereby defining these paths to the control unit as being associated with the same LPAR. See also *logical partition*.

PAV See *parallel access volume*.

PCI See *peripheral component interconnect*.

PDU See *protocol data unit*.

PE See *IBM product engineering*.

peripheral component interconnect (PCI)

An architecture for a system bus and associated protocols that supports attachments of adapter cards to a system backplane.

persistent FlashCopy

A state where a FlashCopy relationship remains indefinitely until the user deletes it. The relationship between the source and target volumes is maintained after a background copy completes.

physical path

A single path through the I/O interconnection fabric that attaches two units. For Copy Services, this is the path from a host adapter on one DS6000 (through cabling and switches) to a host adapter on another DS6000.

pinned data

Data that is held in cache until either an error condition is corrected and it can be moved to disk storage or until the data is discarded by a host command. Pinned data conditions can only occur on an ESS Model 800 during fast-write or dual-copy functions.

point-in-time copy

A FlashCopy option that creates an instantaneous view of original source data at a specific moment in time.

point-to-point connection

A fibre-channel topology that enables the direct interconnection of ports. See also *arbitrated loop* and *switched fabric*.

port A physical connection on a host adapter to the cable that connects the DS6000 to hosts, switches, or another DS6000. The DS6000 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 DS6000 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 DS6000, however, can be configured to indicate that an AS/400 volume is protected or unprotected and give it RAID protection in either case.

protocol data unit (PDU)

A unit of data specified in the protocol of a given layer and consisting of protocol control information for the layer and, possibly, user data for the layer.

pSeries

The product name of an IBM e(logo)server product that emphasizes performance. It is the successor to the RS/6000 family of servers.

pseudo host

A host connection that is not explicitly defined to the DS6000 and that has access to at least one volume that is configured on the DS6000. The FiconNet pseudo host icon represents the FICON 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 DS6000 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 DS6000 automatically reserves spare disk drives when it assigns arrays to a device adapter pair (DA pair). See also *device adapter*, *RAID 10*, and *redundant array of independent disks*.

RAID 10

A type of RAID that optimizes high performance while maintaining fault tolerance for up to two failed disk drives by striping volume data across several disk drives and mirroring the first set of disk drives on an identical set. The DS6000 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 DS6000 that creates a point-in-time data copy without copying an entire volume for each point-in-time copy.

Remote Mirror and Copy

A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also *Global Mirror*, *Metro Mirror* and *Global Copy*.

remote technical assistance information network (RETAIN)

The initial service tracking system for IBM service support, which captures heartbeat and call-home records. See also *support catcher* and *support catcher telephone number*.

REQ/ACK

See *request for acknowledgment and acknowledgment*.

request for acknowledgment and acknowledgment (REQ/ACK)

A cycle of communication between two data transport devices for the purpose of verifying the connection, which starts with a request for acknowledgment from one of the devices and ends with an acknowledgment from the second device. The REQ and ACK signals help to provide uniform timing to support synchronous data transfer between an initiator and a target. The objective of a synchronous data transfer method is to minimize the effect of device and cable delays.

reserved allegiance

For zSeries, a relationship that is created in a control unit between a device and a channel path, or path group, when the device completes a Sense Reserve command. The allegiance causes the control unit to guarantee access (that is, busy status is not presented) to the device. Access is over the set of channel paths that are associated with the allegiance; access is for one or more channel programs until the allegiance ends.

RETAIN

See *remote technical assistance information network*.

S

S/390 IBM enterprise servers based on

Enterprise Systems Architecture/390 (ESA/390). *S/390* is the currently accepted shortened form of the original name *System/390*.

S/390 storage

Storage arrays and logical volumes that are defined as connected to S/390 servers. This term is synonymous with count-key-data storage.

SAID See *system adapter identification number*.

SAM See *sequential access method*.

SAN See *storage area network*.

SBCON

See *Single-Byte Command Code Sets Connection*.

screen The physical surface of a display device upon which information is shown to users.

SCSI See *small computer system interface*.

SCSI device

A disk drive connected to a host through an I/O interface using the SCSI protocol. A SCSI device is either an initiator or a target. See also *initiator* and *small computer system interface*.

SCSI-FCP

Synonym for Fibre Channel Protocol, a protocol used to transport data between an open-systems host and a fibre-channel adapter on an DS6000. See also *Fibre Channel Protocol* and *small computer system interface*.

SCSI host systems

Host systems that are attached to the DS6000 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 DS6000 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 DS6000, for example, in each host bay and in each cluster.

service clearance

The area that is required to open the service covers and to pull out components for servicing.

service information message (SIM)

A message sent by a storage server to service personnel through an zSeries operating system.

service personnel

A generalization referring to individuals or companies authorized to service the

- DS6000. 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 DS6000. This term also refers to a service provider, a service representative, or an IBM service support representative (SSR). An IBM SSR installs the DS6000.
- 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 DS6000 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.
- 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 DS6000 that can replace a failed disk drive. A spare can be predesignated to allow automatic dynamic sparing. Any data preexisting on

a disk drive that is invoked as a spare is destroyed by the dynamic sparing copy process.

spatial reuse

A feature of serial storage architecture that enables a device adapter loop to support many simultaneous read/write operations. See also *serial storage architecture*.

SSID See *subsystem identifier*.

SSR See *service support representative*.

stacked status

For zSeries, the condition when the control unit is in a holding status for the channel, and the last time the control unit attempted to present the status, the channel responded with the stack-status control.

stage operation

The operation of reading data from the physical disk drive into the cache.

staging

To move data from an offline or low-priority device back to an online or higher priority device, usually on demand of the system or on request of the user.

standard volume

A volume that emulates one of several zSeries volume types, including 3390-2, 3390-3, 3390-9, 3390-2 (3380-track mode), or 3390-3 (3380-track mode), by presenting the same number of cylinders and capacity to the host as provided by the native zSeries volume type of the same name.

STI See *self-timed interface*.

storage area network

A network that connects a company's heterogeneous storage resources.

storage capacity

The amount of data that a storage medium can hold; usually expressed in kilobytes, megabytes, or gigabytes.

storage complex

A configuration of one or more storage units that is managed by a management console.

storage device

A physical unit that provides a mechanism to store data on a given

medium such that it can be subsequently retrieved. See also *disk drive module*.

storage extent

The minimum contiguous range of storage on a physical storage device, array, or rank that can be allocated to a local volume

storage image

A partitioning of a storage unit that provides emulation of a storage server with one or 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 DS6000 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 DS6000 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 DS6000 host adapter for use by Copy Services.

System Management Interface Tool (SMIT)

An interface tool of the AIX operating system for installing, maintaining, configuring, and diagnosing tasks.

System Modification Program

A program used to install software and software changes on MVS systems.

T

target A SCSI device that acts as a subordinate to an initiator and consists of a set of one or more logical units, each with an assigned logical unit number (LUN). The logical units on the target are typically I/O devices. A SCSI target is analogous to a zSeries control unit. See also *small computer system interface*.

target device

One of the devices in a dual-copy or remote-copy volume pair that contains a duplicate of the data that is on the source device. Unlike the source device, the target device might only accept a limited subset of data. See also *source device*.

TB See *terabyte*.

TCP/IP

See *Transmission Control Protocol/Internet Protocol*.

terabyte (TB)

- 1) Nominally, 1 000 000 000 000 bytes, which is accurate when speaking of bandwidth and disk storage capacity.
- 2) For DS6000 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 zSeries, the address associated with a device on a given control unit. On OEMI interfaces, the unit address specifies a control unit and device pair on the interface.

unprotected volume

An AS/400 term that indicates that the AS/400 host recognizes the volume as an unprotected device, even though the storage resides on a RAID-formatted array and is, therefore, fault tolerant by definition. The data in an unprotected volume can be mirrored. Also referred to as an *unprotected device*.

upper-layer protocol

The layer of the Internet Protocol (IP) that supports one or more logical protocols (for example, a SCSI-command protocol and an ESA/390 command protocol). Refer to ANSI X3.230-199x.

UTC See *Coordinated Universal Time*.

V**virtual machine facility**

A virtual data processing machine that appears to the user to be for the exclusive use of that user, but whose functions are accomplished by sharing the resources of a shared data processing system. An alternate name for the VM/370 IBM operating system.

vital product data (VPD)

Information that uniquely defines the system, hardware, software, and microcode elements of a processing system.

VM The root name of several IBM operating

systems, such as VM/XA, VM/ESA, VM/CMS, and z/VM. See also *virtual machine facility*.

volume

For zSeries, the information recorded on a single unit of recording medium. Indirectly, it can refer to the unit of recording medium itself. On a nonremovable-medium storage device, the term can also indirectly refer to the storage device associated with the volume. When multiple volumes are stored on a single storage medium transparently to the program, the volumes can be referred to as logical volumes.

volume group

A collection of either physical or logical volumes.

volume label

A unique identifier that a user assigns to a logical volume.

VPD See *vital product data*.

VSE/ESA

An IBM operating system, the letters of which represent virtual storage extended/enterprise systems architecture.

W

weight distribution area

The area that is required to distribute the weight of the storage unit.

worldwide node name (WWNN)

A unique 64-bit identifier for a host that contains a fibre-channel port. See also *worldwide port name*.

worldwide port name (WWPN)

A unique 64-bit identifier associated with a fibre-channel adapter port. It is assigned in an implementation- and protocol-independent manner. See also *worldwide node name*

write hit

A write operation in which the requested data is in the cache.

write penalty

The performance impact of a classical RAID-5 write operation.

WWNN

See *worldwide node name*.

WWPN

See *worldwide port name*.

X

xSeries

The product name of an IBM e(logo)server product that emphasizes industry-standard server scalability and self-managing server technologies. It is the successor to the Netfinity family of servers.

Z

z/Architecture

An IBM architecture for mainframe computers and peripherals. The IBM e(logo)server zSeries family of servers uses the z/Architecture architecture. It is the successor to the S/390 and 9672 family of servers. See also *iSeries*.

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 IBM e(logo)server product line that supports 64-bit real storage.

z/OS Global Mirror

A function of a storage server that assists a control program to maintain a consistent copy of a logical volume on another storage unit. All modifications of the primary logical volume by any attached host are presented in order to a single host. The host then makes these modifications on the secondary logical volume. This function was formerly called *extended remote copy* or *XRC*.

zSeries

An IBM e(logo)server family of servers that emphasizes near-zero downtime.

IBM enterprise servers based on z/Architecture.

zSeries storage

Storage arrays and logical volumes that are defined in the DS6000 as connected to zSeries servers.

Index

Numerics

- 1750-511 and 1750-522 (DS6800) 3
- 1750-EX1 and 1750-EX2 expansion enclosures 5

A

- accessibility 255
- account information 145
- account setup 140
 - unlocking an administrative password 217
- activating licenses 95
- Activation (DSFA), Disk Storage Feature 8
- adapters 27
- address group 182
- AIX
 - CLI installation considerations 105
 - LIBPATH handling 105
- AIX LIBPATH 125
- alert traps 146
- API, DS Open 10
- applying
 - a configuration 175
- arrays
 - creating 168
- arrays, creating (ckd) 191
- arrays, creating (FB) 179
- Assemble tools and equipment 34
- assigning ranks 192

C

- cable problems, determining 228
- calculating hexadecimal values 253
- call home
 - setting up 100
- Call Home
 - account information 145
 - activating 159
 - testing 159
- Call Home support 29
- checking command usage 135
- circuit breakers
 - high-voltage 26
 - low-voltage 26
- ckd
 - creating LCU 193
- CKD configuration error 194
- CKD configuration error, correcting 196
- ckd volume
 - creating arrays 191
- CKD volume rank 192
- ckd volumes, creating 194
- CLI
 - calculating hexadecimal values 253
 - postinstallation 132
- CLI commands on a command line 133
- CLI default configuration profile file 147

- CLI help 136
- CLI history function 135
- CLI interactive mode 135
- CLI, command-line interface 9
- code level 95
- code load 199, 201, 203
- collecting problem determination data 228
- command line presentation 133
- communication requirements, host attachment 27
- concurrent code load 199, 201, 203
- configuration
 - applying 175
 - DS Storage Manager 9
 - offline 22
 - online 22
 - real-time 22
 - reconfiguration 9
 - simulated 22
- configuration error 184
- configuration files
 - creating 91
- Configuration overview 19
- configuration profile file 147
- configuring
 - fibre-channel I/O ports 186, 197
 - I/O ports 102
- configuring I/O port, System i 186, 197
- connect hosts to server enclosure
 - processor cards 61
- connect power cords 63
- connecting expansion enclosures 56
- connecting network cables 55
- connectivity issues, resolving 223
- connectivity, dial-up 29
- console mode installation, DS CLI 113
- contact information 100
- contact information, CLI 142
- contacting IBM 229
- controlling CLI page help 136
- Copy Services
 - disaster recovery 17, 239
 - licensed functions 18
 - overview 12
- copying problem determination data 228
- correcting a CKD configuration error 196
- correcting fixed block configuration error 184
- count key data volumes
 - creating extent pools 189
- creating
 - arrays 168
 - ckd volumes 194
 - extent pools 170
 - extent pools, fixed block 177
 - fixed block volume groups 184
 - fixed block volumes 182
 - host systems 102

- creating (*continued*)
 - open systems volume groups 174
 - open systems volumes 173
 - ranks 169
 - ranks for CKD volumes 192
 - SCSI host port connections 187
 - storage unit 92
 - zSeries volumes 172
 - creating 172
- creating a rank, DS CLI 180
- creating a storage complex 93
- creating arrays (ckd) 191
- creating arrays (FB) 179
- creating configuration files 91
- creating extent pools, CKD 189
- creating LCUs 171
- customer contacts
 - defining 99
- customer name for SMC PE package 233

D

- decimal 253
- default CLI settings 147
- default directories 133
- determining cable problems 228
- determining problems 227
- disaster recovery
 - Copy Services 17, 239
 - failback 240
 - failover 240
 - Global Mirror 245
 - Global Mirror planned scenario 247
 - Global Mirror scenario for unplanned 248
 - two-site solution 245
 - using the fast reverse restore feature 250
 - without using the fast reverse feature 251
- Disk Manager
 - monitoring performance 11
- Disk Storage Feature Activation (DSFA) 8
- display panel
 - removing 37
 - replacing 52
- downloading code 203
- DS CLI
 - console mode installation 113
 - default directories 133
 - first time login 133
 - graphical mode installation 119
 - installation on Linux 113
 - interactive mode 135
 - persistent setup, OpenVMS 131
 - script mode 135
 - silent mode installation 109
 - silent mode installation, DS CLI 109
 - single-shot mode 134

- DS CLI command format 133
- DS CLI installation
 - AIX considerations 105
 - default directories 105
 - HP Tru64 considerations 105
 - mounting the installation CD 105
 - preparing for 105
 - System i considerations 105
- DS command-line interface 9
- DS Open API 10
- DS storage management
 - Windows uninstall 209
- DS Storage Manager 9
- DS6000
 - major features 2
- DS6000 interfaces 9
- DS6000 storage management
 - Windows
 - unattended (silent) mode
 - install 82
 - unattended (silent) mode
 - upgrade 208
- DS6000 Storage Manager
 - Windows graphical install 76
 - Windows graphical upgrade 204
- DS6800
 - downloading code 203
 - upgrading code 199, 201, 203
 - verifying code level 95
- DSFA, Disk Storage Feature Activation 8, 95

E

- e-mail 100
- electrostatic discharge procedures 35
- email 100
- enclosure connectivity, verifying 225
- error messages
 - generic SNMP 217
 - product-specific 219
 - storage unit 221
- ESD procedures 35
- Ethernet cable 55
- exception events 146
- expansion enclosures
 - connecting 56
- express configuration
 - overview 22
- express Configure
 - iSeries volumes 166
 - open systems volumes 163
 - zSeries volumes 165
- extent pool names 177
- extent pools
 - creating 170
- extent pools, System i 177

F

- failover
 - site switch 240, 243
- failover and fallback 17, 239
- fast reverse
 - Global Mirror 250

- fast reverse restore
 - advantages of using 250
- FATA disk drives 2
- Feature Activation (DSFA), IBM Disk Storage 95
- feature licensing 8
- Fibre Channel Protocol 12
- fibrec-channel
 - host attachment 27
- fibrec-channel ATA (FATA) disk drives 2
- fibrec-channel I/O ports
 - configuring 186, 197
 - fc-al 186, 197
 - ficon 186, 197
 - scsi-fcp 186, 197
- FICON 12
- field updates 1
- field upgrade 56, 69
- firmware
 - applying an update 199
 - checking for updates 201
 - downloading an update 203
 - installing an update 203
 - verifying level 95
- first time login 133
- fixed block configuration error, correcting 184
- fixed block volume
 - creating arrays 179
- fixed block volume groups, creating 184
- fixed block volumes
 - creating extent pools 177
- fixed block volumes, creating 182
- following a light path for unguided service 201
- following the problem log for guided service 202
- front display panel 37, 52
 - removing 37
 - replacing 52
 - service effects 37
- FTP transfer failed, pe package 234
- FTP transfer failed, statesave file 236
- FTP, problem determination process 230
- functions
 - Copy Services 12

G

- gateway 66
- Global Mirror 17, 239
 - disaster recovery for unplanned 248
 - disaster recovery without fast reverse 251
 - fast reverse restore 250
 - planned failover and fallback scenario 247
 - setting up 245
 - unplanned failover and fallback 248
- graphical mode installation, DS CLI 119
- guided service 202

H

- help, CLI 136
- hex 253

- hexadecimal 253
- high-voltage
 - installations 26
- host connectivity, verifying 226
- host systems
 - adapter types 12
 - communication requirements 27
 - creating 102

I

- I/O ports
 - configuring 102
- i5/OS
 - command format 134
 - removing DS CLI directly 215
 - silent mode CLI installation 109
- IBM Disk Storage Feature Activation (DSFA) 95
- IBMDSCCLI\$STARTUP.COM
 - procedure 131
- implementation, RAID 7
- Information Center 1
- informational messages 220
- input frequencies 26
- install
 - fiber optic cable 56
 - front display panel 52
 - SFP 56
- installation data 29
- installation overview 2
- installing server enclosure in the rack 44
- installing the CLI 132
- installing the DS6000 2
- Installing the host systems and I/O adapters 53
- installing the support rails 42
- interactive mode
 - history function 135
 - setoutput command 135
- interenclosure connectivity, verifying 225
- interfaces, DS6000 9
- Internet Explorer 74
- IP address 66
- iSeries volumes
 - creating 166

J

- JVM not found error 125

K

- keyboards
 - accessibility features 255

L

- LaunchPad, Windows install 76
- LCU creation 193
- LCUs
 - creating 171
- LED status 65

- legal
 - terms and conditions 258
- licensed functions 8, 95
 - Copy Services 18
- licenses
 - Disk Storage Feature Activation (DSFA) 95
 - function authorization documents 95
- light path
 - analysis 227
 - performing a service action 201
- light path analysis 227
- light path service 201
- line cords
 - power connector 26
- Linux
 - DS CLI installation, console mode 113
- logical control unit creation 193
- logically removing a physical resource 222
- low-voltage
 - installations 26

M

- machine signature 93
- maintenance information 1
- management console 74
 - defining peer 97
- Management Console
 - defining multiple 97
- model 93
- Models 1750-511 and 1750-522 3
- Models 1750-EX1 and 1750-EX2 5
- modem 69, 160
- mounting the installation CD 105
- multiple IP addresses 68
- multiple Management Consoles
 - defining 97

N

- near-line 2
- Netscape Navigator 74
- network mask 66
- network settings 66
- nodes 27
- nonconcurrent code load 199, 201, 203
- notifications
 - configuring 100
- Novell DS CLI graphical installation 119

O

- obtaining activation codes 95
- offline configuration overview 22
- online configuration overview 22
- open systems volume groups
 - creating 174
- open systems volumes
 - creating 163, 173
- OpenVMS
 - ECOs for installation 108
 - IBMD\$CLI\$STARTUP.COM procedure 131

- OpenVMS (*continued*)
 - installing CLI 126
 - mounting the installation CD 126
 - persistent setup 131
 - Polycenter software installation utility 108
 - preparing for CLI installation 108
- OS/400 DS CLI graphical installation 119
- overview
 - Copy Services 12
 - offline configuration 22
 - online configuration 22
 - real-time configuration 22
 - simulated configuration 22

P

- panels, verifying correct operation 226
- peer management console
 - defining 97
- performance features 6
- performance gathering
 - Disk Manager 11
- performing
 - site switch
 - failover 240, 243
- performing a light path analysis 227
- persistent setup 131
- planning
 - installation 27
 - licensed functions 8
 - power connector 26
 - SAN requirements 28
- Polycenter software installation utility 108
- positioning the rack 35
- postinstallation
 - CLI 132
- postinstallation tasks 132
- power
 - outlet requirements 25
- power connector
 - requirements 26
- power on issues, resolving 224
- powering off 238
- powering on 64
- preinstallation planning and verification 25
- preparing
 - installation information 29
- Preparing your system for installation 25
- problem determination datacollecting
 - problem determination data sending 228
- problem determination process, FTP 230
- problem log service 202
- problem states 218
- product-specific error messages 219

R

- RAID
 - implementation 7
 - RAID 10 overview 8

- RAID (*continued*)
 - RAID 5 overview 7
 - RAID overview 7
- rank configuration 180, 192
- rank group 1 or 1 192
- rank, create with DS CLI 180
- ranks
 - creating 169
- real-time configuration overview 22
- recovering 245
 - with and without the fast reverse restore feature 250
- remote support 29, 69
 - activating 160
 - contact information 142
 - VPN connection
 - troubleshooting 160
- Remove the power supply 38
- removing DS CLI, i5/OS 215
- removing logical resources after physical removal 222
- Removing the disk drive module 38
- removing the front display panel 37
- removing the hardware resources 35
- repair
 - front display panel 37, 52
- Replacing the disk drive module 47
- replacing the front display panel 52
- replacing the resources 47
- Replication
 - copy services functions 11
- report formats
 - controlling 135
- requirements
 - floor-load 25
 - host attachment communication 27
 - input voltage 26
 - management console 74
 - power connectors 26
 - power outlets 25
 - service clearance 25
- resolving IP address connectivity issues 223
- resolving power on issues 224
- resources
 - front display panel operation 226
 - front display panel, service effects 37
 - light path service 201
 - problem log service 202
 - rear operator panel operation 226
- routing the cables 53

S

- S/390 host
 - service information messages 218
- S/390 messages 146
- safety
 - ESD 35
 - power outlets 25
- script mode 135
- SCSI host port connections
 - creating 187
- secure shell connection 145
- security recovery utility 97
- sending a pe package 234
- sending a statesave file 236

- sending problem determination data 228
- Serial cable 55
- serial number 93
- serial number, obtaining 138
- service
 - through the light path 201
 - through the problem log 202
- service clearance requirements 25
- service effects for the front display panel 37
- service information messages
 - zSeries and S/390 host 218
- Service information messages 146
- Service Information Messages 100
- setting up
 - Global Mirror 245
- setup, account 140
- signature 93
- SIM 100
- SIM messages 146
- simulated configuration overview 22
- single-shot mode
 - command format, i5/OS 134
 - command format, Windows 134
- SMC PE package 233
- SMTP 100
- SNMP 100
- SNMP notifications 146
- SNMP trap
 - messages 217
- storage image ID, obtaining 138
- storage plex
 - configuring iSeries volumes
 - quickly 166
 - configuring open systems volumes
 - quickly 163
 - configuring zSeries volumes
 - quickly 165
- storage unit
 - assigning to a storage complex 94
 - creating 92
 - downloading code 203
 - error messages 221
 - installing code upgrades 203
 - reviewing network settings 99
 - specifying day and time 98
 - upgrading code 199, 201
 - verifying code level 95
- support operations
 - SNMP notifications 146
- support options, CLI
 - Call Home activation 159
 - e-mail notifications 144
 - SMTP 144
- support rails
 - install 42
- system compatibility 74
- System i
 - bypassing uninstaller 215
 - CLI installation considerations 105
 - configuring I/O port 186, 197
 - extent pool considerations 177
 - host connections 187
 - interactive mode 135
 - postinstallation tasks 132
 - removing CLI, direct method 215
 - removing CLI, remote method 216

- System i (*continued*)
 - script mode 135
 - single-shot mode 134

T

- timeout connection 147
- TotalStorage Productivity Center
 - copy services 11
 - for Disk 11
 - Replication 11
- Trademarks 259
- troubleshooting
 - cables 228
 - DS CLI commands, FTP 230
 - generic SNMP trap messages 217
 - informational messages 220
 - normal operation problems 220
 - problem states 218
 - product-specific error messages 219
 - service information messages 218
 - storage unit error messages 221
- Troubleshooting
 - DS CLI commands, no ftp 237

U

- unattended (silent) mode
 - installation 109
- unattended (silent) mode installation, DS CLI 109
- Unconfigure modem 160
- unconfiguring 160
- unconfiguring modem 160
- unguided service 201
- uninstall DS storage management
 - Windows 209
- uninstall.dat 216
- uninstall.jar 216
- UNIX DS CLI graphical installation 119
- upgrading code 199, 201, 203
- user administration
 - unlocking an administrative password 217
- using
 - fast reverse restore 250

V

- verify status through LEDs 65
- verify the ship group 33
- verifying 95
- verifying code level 95
- verifying correct operation of the panels 226
- verifying host connectivity 226
- verifying interenclosure connectivity 225
- voltage requirements, input 26
- VPN connection 69
- VPN connectiontroubleshooting for remote support 160

W

- Windows
 - DS Storage Manager upgrade, graphical 204
 - graphical install, DS6000 Storage Manager 76
 - install, LaunchPad 76
 - uninstall, DS storage management 209
- Windows DS CLI graphical installation 119
- Windows unattended (silent) mode install 82
- Windows unattended (silent) mode upgrade 208

Z

- zones
 - WWPN 12
- zSeries host
 - service information messages 218
- zSeries system messages 146
- zSeries volumes
 - creating quickly 165

Readers' comments — we would like to hear from you

IBM System Storage DS6000
Installation, Troubleshooting, and Recovery Guide
Version 1 Release 2

Publication No. GC26-7925-00

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.

Submit your comments using one of these channels:

- 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



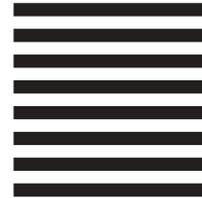
NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM
Information Development
Department 61D
9032 South Rita Road
Tucson, Arizona 85775-4401



Fold and Tape

Please do not staple

Fold and Tape



Part Number: 22R6312

Printed in USA

GC26-7925-00



(1P) P/N: 22R6312



Spine information:



IBM System Storage DS6000

Installation, Troubleshooting, and Recovery Guide

Version 1
Release 2