IBM TotalStorage DS6000

Installation, Troubleshooting, and Recovery Guide

GC26-7678-05

IBM TotalStorage DS6000

Installation, Troubleshooting, and Recovery Guide

GC26-7678-05

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

Sixth Edition (June 2005)

Contents

Tables

About this guide

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

Who should use this guide

This publication is for anyone who is planning to install, troubleshoot, or recover a IBM® TotalStorage® DS6000®.

Safety and environmental notices

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

Safety notices

Use this process to find information about safety notices.

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

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

   DANGER

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

   1000

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

   1001

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

Environmental notices

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

Product recycling

This unit contains recyclable materials.

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

Disposing of products

This topic contains information about how to dispose of products.

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

Conventions used in this guide

The following typefaces are used to show emphasis:

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

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

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

Related information

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

- The publications that make up the IBM® TotalStorage™ 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 xvi for information about how to order publications in the IBM TotalStorage DS6000 series publication library. See "How to send your comments" on page xvii 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-8803. These publications are also available as PDF files by clicking on the **Documentation link** on the following Web site:

http://www-1.ibm.com/servers/storage/support/disk/ds6800/index.html

See "Ordering IBM publications" on page xvi for information about ordering these and other IBM publications.

Title	Description	Order Number
IBM® TotalStorage® DS: Command-Line Interface User's Guide	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.)

Title	Description	Order Number
IBM TotalStorage DS6000: Host Systems Attachment Guide	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.)
IBM TotalStorage DS6000: Introduction and Planning Guide	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
IBM TotalStorage Multipath Subsystem Device Driver User's Guide	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 redundant paths for shared logical unit numbers. SDD permits applications to run without interruption when path errors occur. It balances the workload across paths, and it transparently integrates with applications.	SC30-4096
IBM TotalStorage DS Application Programming Interface Reference	This publication provides reference information for the IBM TotalStorage DS application programming interface (API) and provides instructions for installing the Common Information Model Agent, which implements the API.	GC35-0493
IBM TotalStorage DS6000 Messages Reference	This publication provides explanations of error, information, and warning messages that are issued from the DS6000 user interfaces.	GC26-7682
IBM TotalStorage DS6000 Installation, Troubleshooting, and Recovery Guide	This publication provides reference information for installing and troubleshooting the DS6000. It also discusses disaster recovery using Copy Services.	GC26-7678
IBM TotalStorage DS6000 Quick Start Card	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/ds6800/index.html		

Other IBM publications

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

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

Title	Description	Order Number
Data-copy services		
z/OS DFSMS Advanced Copy Services	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

Title	Description	Order Number
IBM Enterprise Storage Server	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
Implementing Copy Services 0n S/390	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
IBM TotalStorage ESS Implementing Copy Services in an Open Environment	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on UNIX, Windows NT®, Windows 2000, Sun Solaris, HP-UX, Tru64, OpenVMS, and iSeries host systems. The Copy Services functions that are described include peer-to-peer remote copy and FlashCopy. This publication describes the functions and shows you how to implement them into your environment. It also shows you how to implement these functions in a high-availability cluster multiprocessing environment.	SG24-5757
Fibre channel		
Fibre Channel Connection (FICON) I/O Interface: Physical Layer	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
Fibre Transport Services (FTS): Physical and Configuration Planning Guide	This publication provides information about fibre-optic and ESCON-trunking systems.	GA22-7234
IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08.	SC26-7350
IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide	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
IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide	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
IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide	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
Implementing Fibre Channel Attachment on the ESS	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		

Title	Description	Order Number
ESS Solutions for Open Systems Storage: Compaq AlphaServer, HP, and Sun	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
IBM TotalStorage ESS Implementing Copy Services in an Open Environment	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
Implementing Fibre Channel Attachment on the ESS	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		
Device Support Facilities: User's Guide and Reference	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
z/OS Advanced Copy Services	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
DFSMS/MVS V1: Remote Copy Guide and Reference	This publication provides guidelines for using remote copy functions with S/390 and zSeries hosts.	SC35-0169
Fibre Transport Services (FTS): Physical and Configuration Planning Guide	This publication provides information about fibre-optic and ESCON-trunking systems.	GA22-7234
Implementing ESS Copy Services on S/390	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on an Enterprise Storage Server that is attached to an S/390 or zSeries host system. Copy Services functions include peer-to-peer remote copy, extended remote copy, FlashCopy, and concurrent copy. This publication describes the functions, prerequisites, and corequisites and describes how to implement each function into your environment.	SG24-5680

Title	Description	Order Number
ES/9000, ES/3090: IOCP User Guide Volume A04	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
IOCP User's Guide, IBM e(logo)server zSeries 800 and 900	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
IOCP User's Guide, IBM e(logo)server zSeries	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
S/390: Input/Output Configuration Program User's Guide and ESCON Channel-to-Channel Reference	This publication describes the Input/Output Configuration Program that supports ESCON architecture and the ESCON multiple image facility.	GC38-0401
IBM z/OS Hardware Configuration Definition User's Guide	This guide provides conceptual and procedural information to help you use the z/OS Hardware Configuration Definition (HCD) application. It also explains: • 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
OS/390: Hardware Configuration Definition User's Guide	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
OS/390 V2R10.0: MVS System Messages Volume 1 (ABA - ASA)	This publication lists OS/390 MVS™ system messages ABA to ASA.	GC28-1784
Using IBM 3390 Direct Access Storage in a VM Environment	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
Using IBM 3390 Direct Access Storage in a VSE Environment	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
Using IBM 3390 Direct Access Storage in an MVS Environment	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
z/Architecture Principles of Operation	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

Title	Description	Order Number
SAN		
IBM OS/390 Hardware Configuration Definition User's Guide	This guide explains how to use the Hardware Configuration Data application to perform the following tasks: • 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
IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08.	SC26-7350
IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide	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
IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide	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
IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide	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
Implementing Fibre Channel Attachment on the ESS	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		
Device Support Facilities: User's Guide and Reference	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
IBM TotalStorage Solutions Handbook	This handbook, from the IBM International Technical Support Organization, helps you understand what makes up enterprise storage management. The concepts include the key technologies that you must know and the IBM subsystems, software, and solutions that are available today. It also provides guidelines for implementing various enterprise storage administration tasks so that you can establish your own enterprise storage management environment.	SG24-5250

Ordering IBM publications

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

IBM publications center:

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

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

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

Publications notification system:

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

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

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

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

Web sites

The following Web sites provide information about the IBM TotalStorage 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.html
DS6000 series publications	http://www-1.ibm.com/servers/storage/support/disk/ds6800/index.html Click **Documentation**.
FlashCopy for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
Host system models, operating systems, and adapters that the storage unit supports	http://www.ibm.com/servers/storage/disk/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 TotalStorage DS6000 series	http://www-1.ibm.com/servers/storage/disk/ds6000

Type of Storage Information	Web Site
IBM version of the Java (JRE) that is often required for IBM products	http://www-106.ibm.com/developerworks/java/jdk/
Multiple Device Manager (MDM)	http://www.ibm.com/servers/storage/support/ Click **Storage Virtualization**.
Remote Mirror and Copy (formerly PPRC) for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
SAN fibre channel switches	http://www.ibm.com/storage/fcswitch/
Storage Area Network Gateway and Router	http://www-1.ibm.com/servers/storage/support/san/index.html
Subsystem Device Driver (SDD)	http://www-1.ibm.com/servers/storage/support/software/sdd.html
z/OS Global Mirror (formerly XRC) for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this information or any other 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 GC26-7678-05 IBM TotalStorage D6000 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.

Changed Information

Updates have been made to the following sections:

- “DS6000 Configuration overview” on page 63
- “Creating real-time configurations” on page 63
- “Creating simulated configurations” on page 97

Chapter 1. DS6000 series installation overview

The topics in this publication provide installation, troubleshooting, and recovery information related to your DS6000. Topics covered include installing, upgrading, and removing the storage manager and the command-line interface.

To install the DS6000, follow the steps below. Supporting information is contained within this publication.

	Installation Steps
1	Prepare the physical site for hardware installation.
2	Install the server enclosure and storage enclosures in the rack.
3	Route the cables.
4	Check power and safety.
5	Install the PC you will use as your management console.
6	Begin the DS Storage Manager installation.
7	Complete the DS Storage Manager installation.
8	Complete the DS Storage Manager post-installation tasks.
9	Complete the DS Command-Line Interface (CLI) installation.
10	Complete the DS CLI postinstallation instruction
11	Enable remote support functions.

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

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

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

Verify 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:

   - These items are contained within the enclosure:

  - Two 511 processor cards (server enclosure only) or two EX1 processor cards (storage enclosure only)
  - Two power supplies/fan assemblies
  - Two battery backup units (server enclosure only) or two battery blanks (storage 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 2.8 meter inline power cords
  - One Ethernet crossover cable (server enclosure only)
  - One serial conversion cable (server enclosure only)
  - Two 25 meter Ethernet cables (server enclosure only)
- Software, including:
  - Microcode CD
  - CLI CD
  - SMC software CD (SDD is included)
- License Machine Code Agreement
- Statement of Limited Warranty
- Code Reference Flyer
- Electrostatic discharge (ESD) wrist strap

3. Ensure that the items listed in your packing slip match what is in the box to include any optional items you ordered. For example, if you ordered Fibre channel cables, SFPs, or optional power cords, ensure you received them.
4. If any items are missing or damaged, contact your IBM customer support before proceeding.

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, to place the hardware resources.
2. When you are ready, begin to prepare the rack.

ESD procedures

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.

- 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 instructed to do so, switch off the machine power before you remove ESD-sensitive parts.
- Just before touching 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.
- Never touch or work on any electronic circuits without wearing the ESD wrist strap.
- 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.

Chapter 2. Installing storage and server enclosures in a rack

This section provides information on how to position your rack and install the necessary components.

Note: You have the option to install your software before, or in parallel to, installing the hardware. See Chapter 5, “Installing the DS6000 Storage Manager,” on page 39.

Position 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 storage 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 will allow 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.

Remove the hardware resources

This section provides resource removal instructions.

The purpose of resource removal is to minimize the weight of the DS6000 before you install it in the rack. However, if you have three people available to lift and install the DS6000 in a rack, you might not find it necessary to remove the resources before installation. If this is the case, you can skip the resource removal instructions provided in this section.

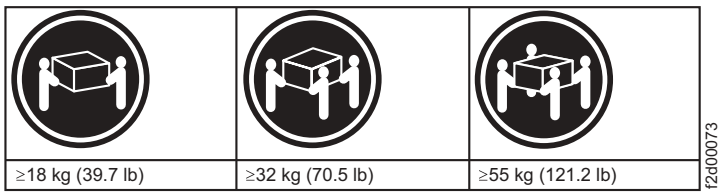

Note: The fully populated enclosure weighs in excess of 109 lbs (49.5 kg). You should have at least three people available to lift and install a fully populated DS6000 in the rack. If there are less than three people to lift the DS6000 from the box and into the rack, you will need to remove the cardboard sides of the shipping box from around the enclosure and remove components from the enclosure before lifting it, as shown below.

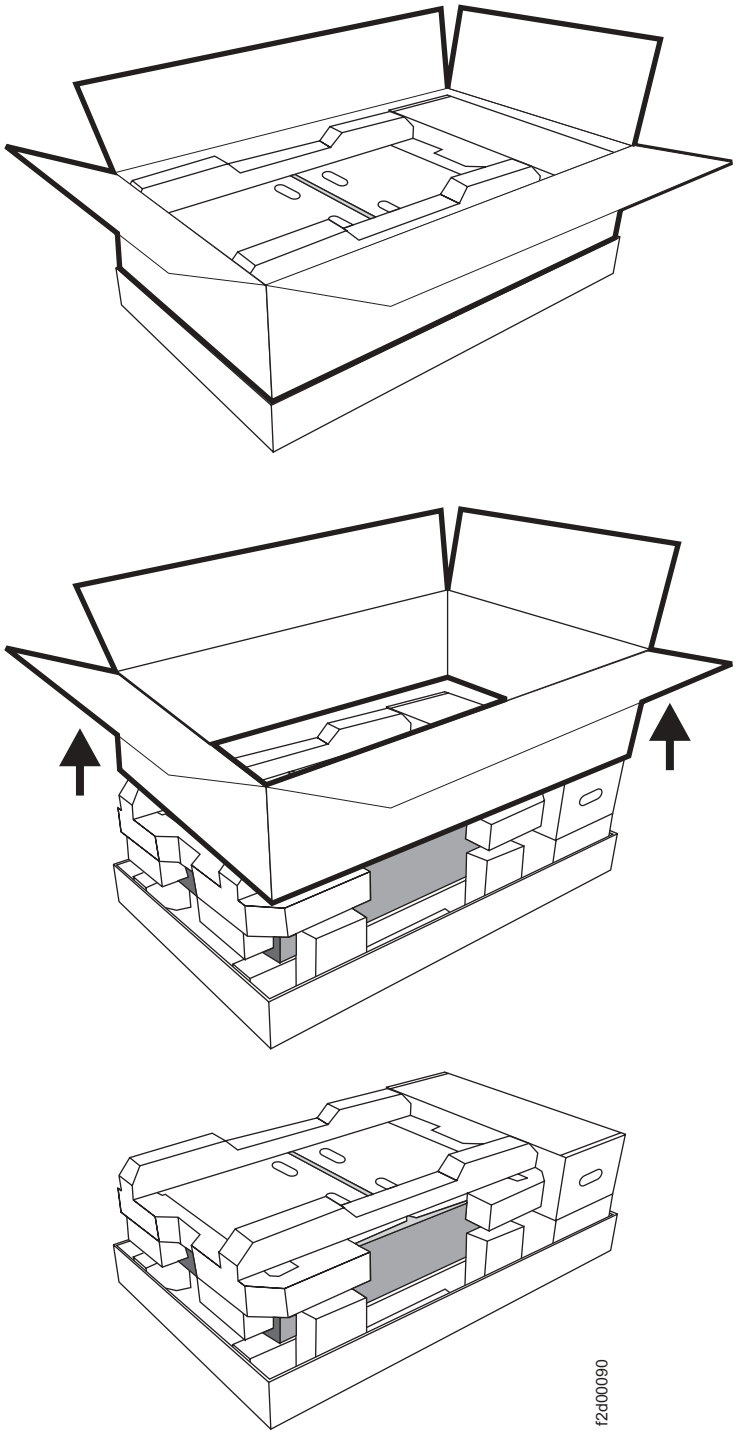

1. Make sure to follow ESD procedures by using an anti-static wrist strap. You will also need a cart or level surface to hold the resources.
2. Remove all the resources in the DS6000 storage and server enclosures if needed to minimize the weight.

Removing the front display panel

Follow these steps to remove the front display panel from the storage enclosure, either when the system is powered off or powered on.

Use approved ESD procedures to prevent damage.

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

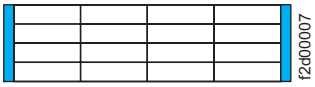

1. Press the blue release button. This releases the handle.
2. Pull the handle up. This action partially ejects the bottom of the display panel.
3. Use the handle to pull the bottom of the panel up and release the top portion of the panel. This completely removes the display panel.

These images show the steps of a front display panel removal procedure.

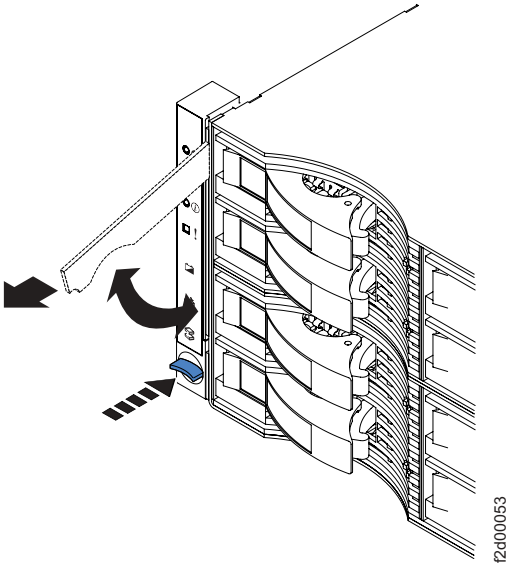

You must replace the front display panel with another front display panel that is in working condition.

Removing the disk drive module

Use approved ESD procedures to prevent damage.

To lessen the weight of the server or storage enclosure during installation, use the following procedure to remove the disk drive modules. The disk drive modules are located in rows and columns on the front of the storage unit.

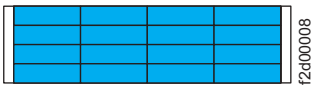

Always handle a defective disk drive module carefully. Damage to defective disk drive modules can have a negative effect on failure analysis test results and warranty recovery.

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 could result in physical damage to the drive or the drive component.

These images show the steps of a disk drive module removal procedure.

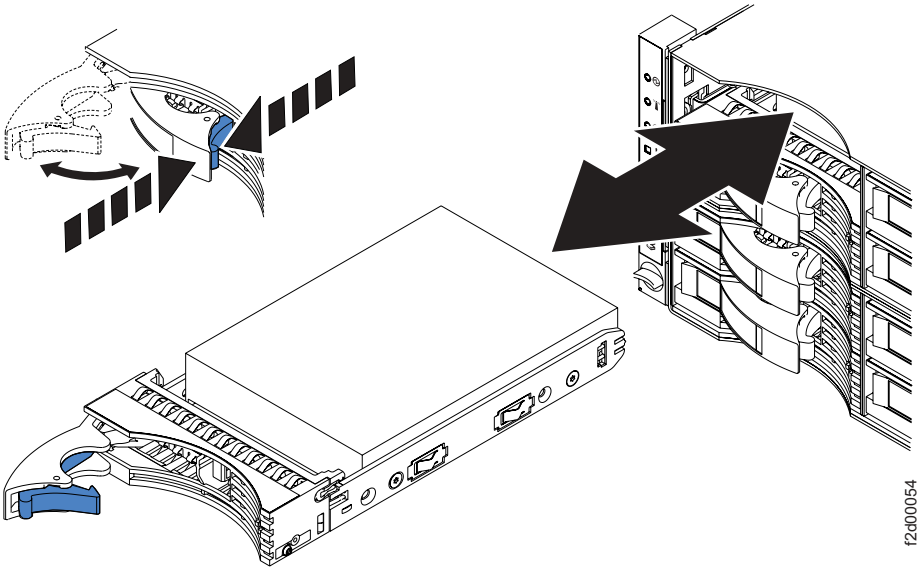

All disk drive module slots should be filled. You can insert a disk drive module blank to prevent overheating the storage unit.

Remove the power supply

Use approved ESD procedures to prevent damage.

To lessen the weight of the server or storage enclosure during installation, use the following procedure to remove the power supplies. The power supplies are located on right and left sides of the rear of the enclosure.

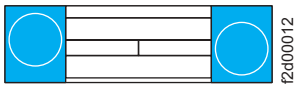

1. Press the blue release button. This releases the handle.
2. Pull the handle out and towards the middle of the enclosure. This action partially ejects the power supply from the slot.
3. Use the handle to pull the power supply partially out of the slot.
4. Grip the power supply with both hands to pull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

These images show a power supply removal procedure.

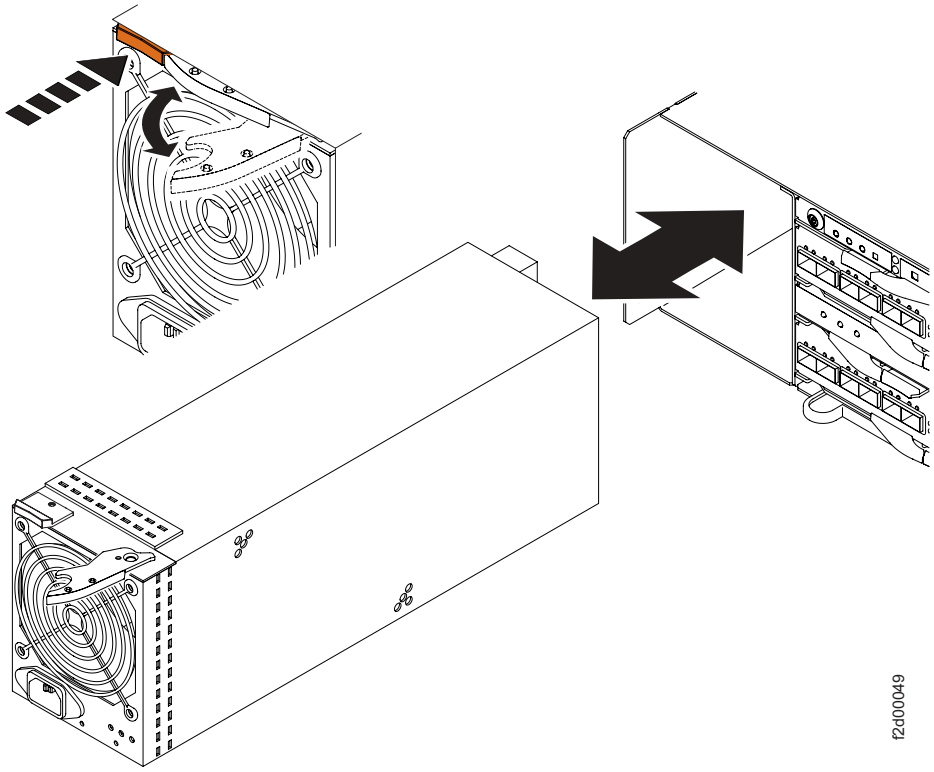

Removing the rear display panel

Use approved ESD procedures to prevent damage.

To lessen the weight of the server or storage enclosure during installation, use the following procedure to remove the rear display panel. The rear operator panel is located at the rear of the storage enclosure, on the top of the resource section.

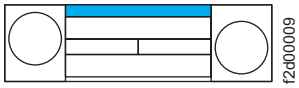

1. Press the orange release button. This releases the handle.
2. Pull the handle out and to the right. This action partially ejects the operator panel from the slot.
3. Use the handle to pull the rear operator panel partially from the slot.
4. Grip the rear operator panel with both hands to pull the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

These images show the steps of a rear operator panel removal procedure.

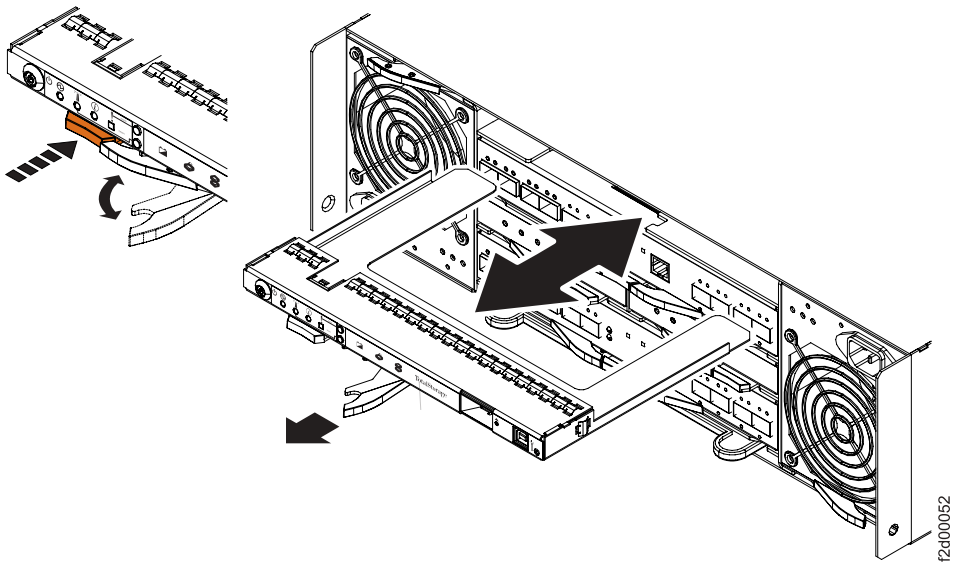

Removing the battery backup unit

Use approved ESD procedures to prevent damage.

To lessen the weight of the server or storage enclosure during installation, use the following procedure to remove the battery backup units. The battery backup unit is located horizontally in the middle of the storage enclosure.

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 the steps of a battery backup unit removal procedure.

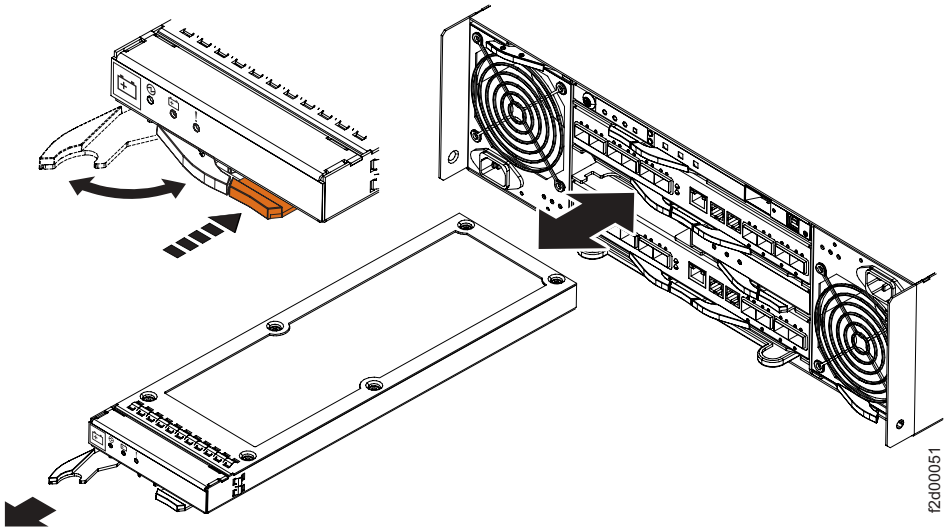

Removing the processor card

Use approved ESD procedures to prevent damage.

To lessen the weight of the server or storage enclosure during installation, use the following procedure to remove the processor cards. The processor cards are located in the rear of the storage 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.

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 the steps of a processor card removal procedure.

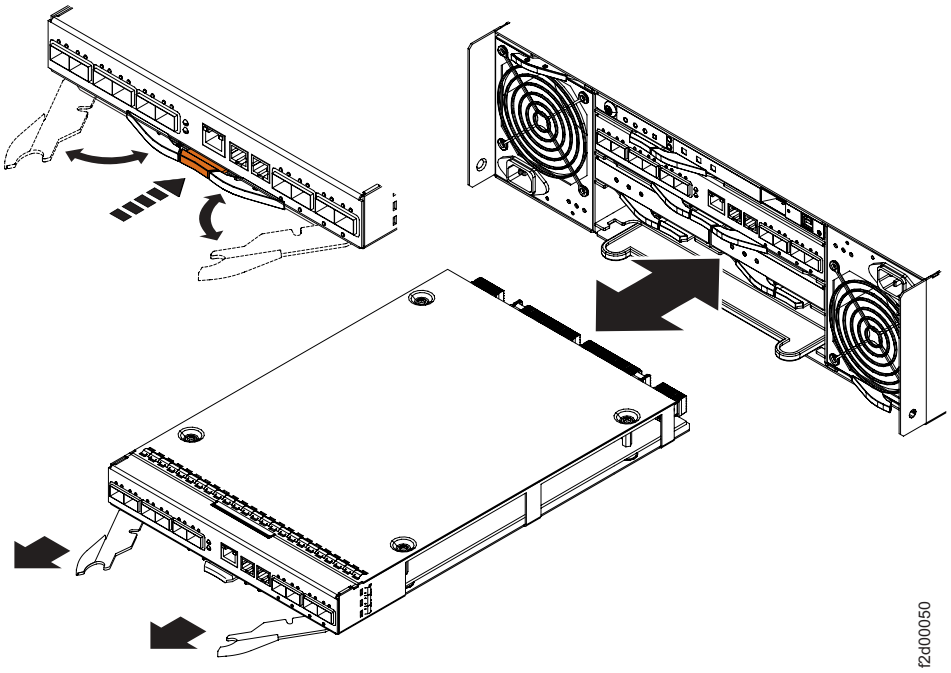
f2d00050

Install the support rails

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

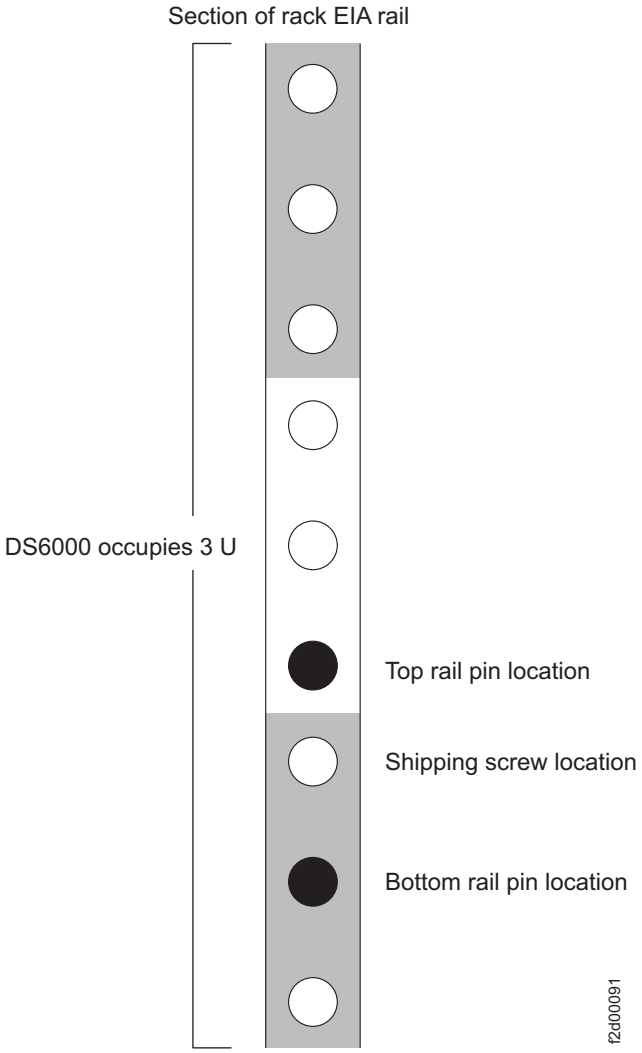

The DS6000 requires one of the following supported IBM racks:

- 2101-200
- 7014
- 9308 (Netbay)

or a compatible Electronic Industries Association (EIA) 310-D Section 1 19-inch rack cabinet. 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 storage enclosure.

Follow these steps to install the support rails:

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.

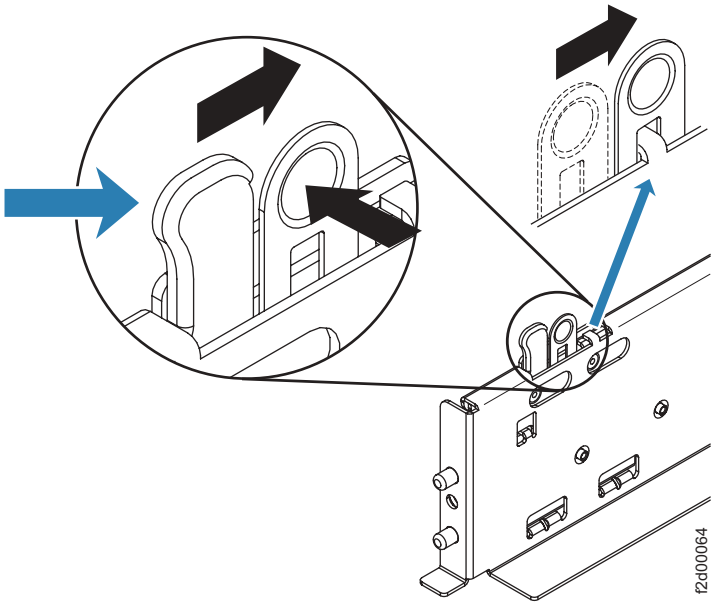

2. With rail pins open, fit the mounting flange around the rack cabinet rail. Align the rail pins with the holes in the rack cabinet rail and release the rail pins.

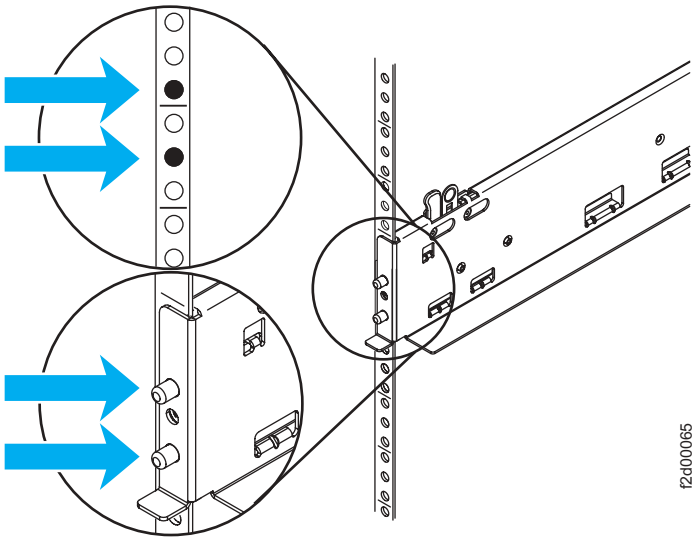

3. Repeat steps two and three with the other end of the rail, aligning with the rear mounting flange.
4. Repeat steps two through four with the left slide rail.

Install 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 storage 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 should have four people available 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 hole 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.

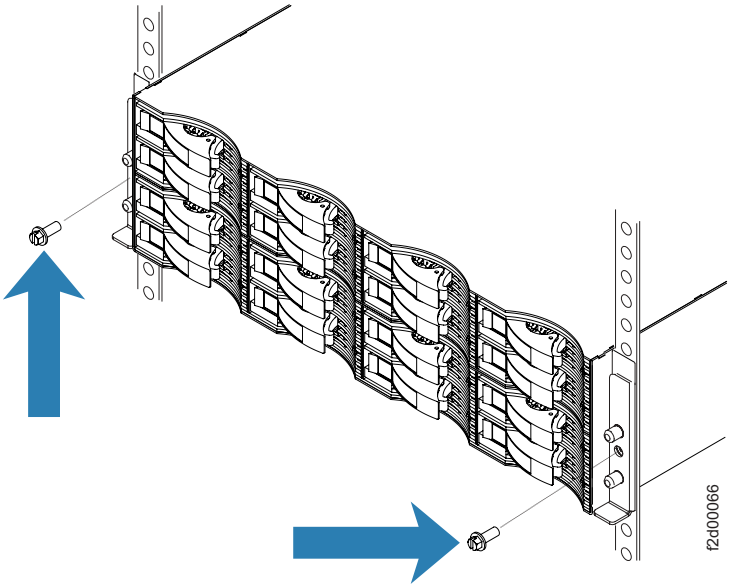

   c. Ensure the rear screw mounting bracket is pushed up against the rear of the chassis. Remove the M5 countersunk screw from the bracket and insert into the hole in the rear of the chassis and tighten. Repeat for the opposite side.

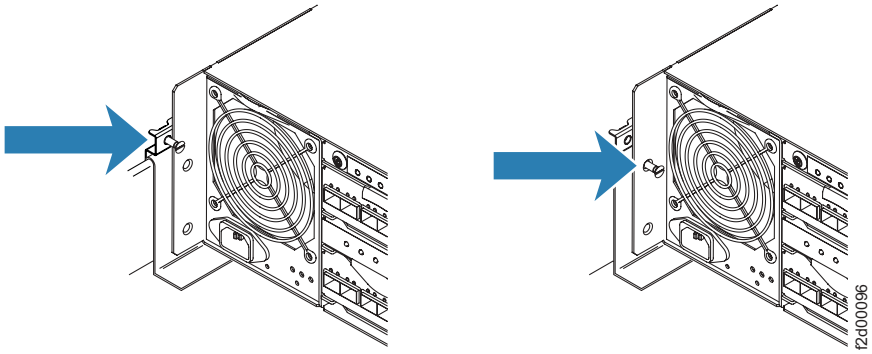

   d. Insert an M5 hex screw into the mounting hole between the two pins on both sides of the rear of the chassis. Tighten the screws to secure the rear of the support rails to both of the rear cabinet rails.

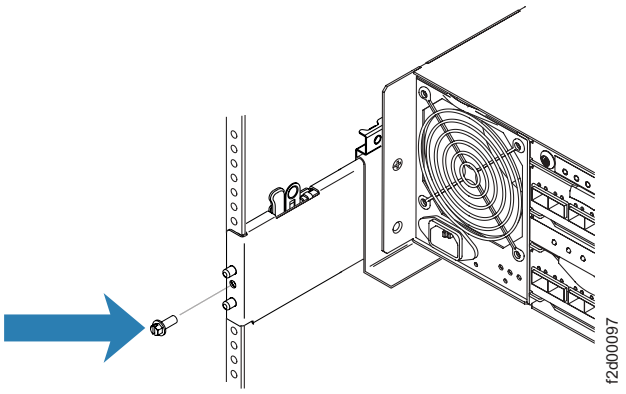

Install storage enclosures in the rack

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

1. Place the storage 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 four people available to lift the unit into the rack.

   a. Move the storage enclosure to the front of the rack-mount cabinet.
   b. Slide the unit into the rack.
2. Secure the storage 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 storage enclosure to both of the front cabinet rails.

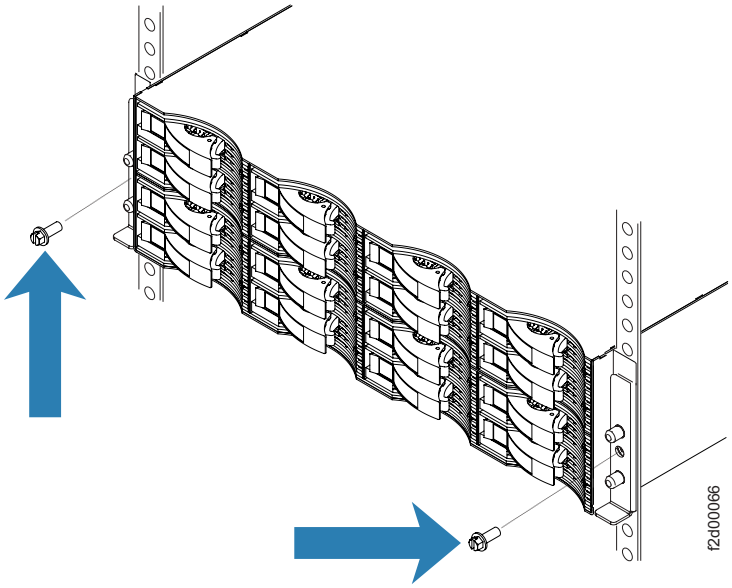

   c. Ensure the rear screw mounting bracket is pushed up against the rear of the chassis. Remove the M5 countersunk screw from the bracket and insert into the hole in the rear of the chassis and tighten. Repeat for the opposite side.

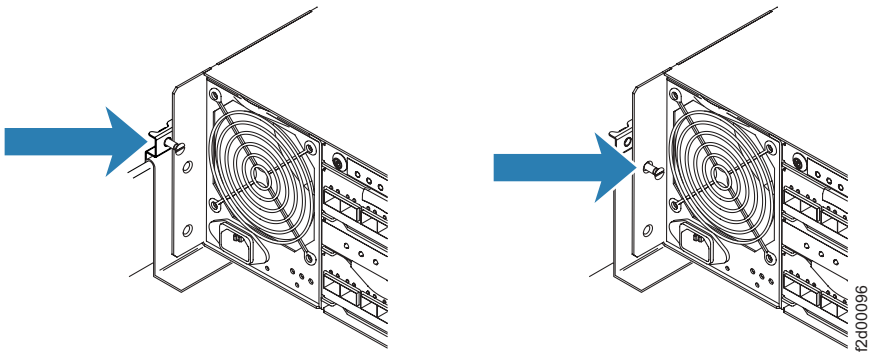

d. Insert an M5 hex screw into the mounting hole between the two pins on both sides of the rear of the chassis. Tighten the screws to secure the rear of the support rails to both of the rear cabinet rails.

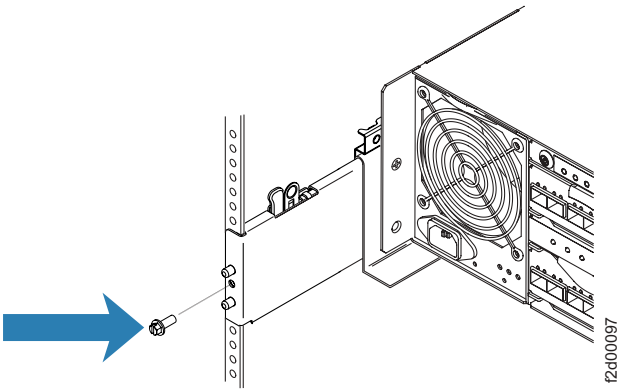

Replace the hardware resources

If you did not remove the resources prior to installation, you can skip this section.

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

Replacing the disk drive module

Use approved ESD procedures to prevent damage.

If you removed the disk drive modules to lessen the weight of the enclosure, use the following instructions to replace the disk drive modules. The disk drive modules are located in rows and columns on the front of the storage unit.

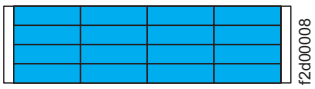

1. Remove the factory-sealed wrapping from the new disk drive module only when you are ready to install it.
2. Before installing the disk drive module, open the disk drive module handle by pressing the blue latch and pulling the handle open.
3. Align the disk drive module with the groove on the disk 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.

4. Push the disk drive module handle to the right until it is latched closed.
5. Verify that the front of the new disk drive module is aligned with the other disk drive modules.

   The storage unit will automatically begin the process to bring the DDM online and rebuild the array.

These images show the steps of a disk drive module replacement procedure.

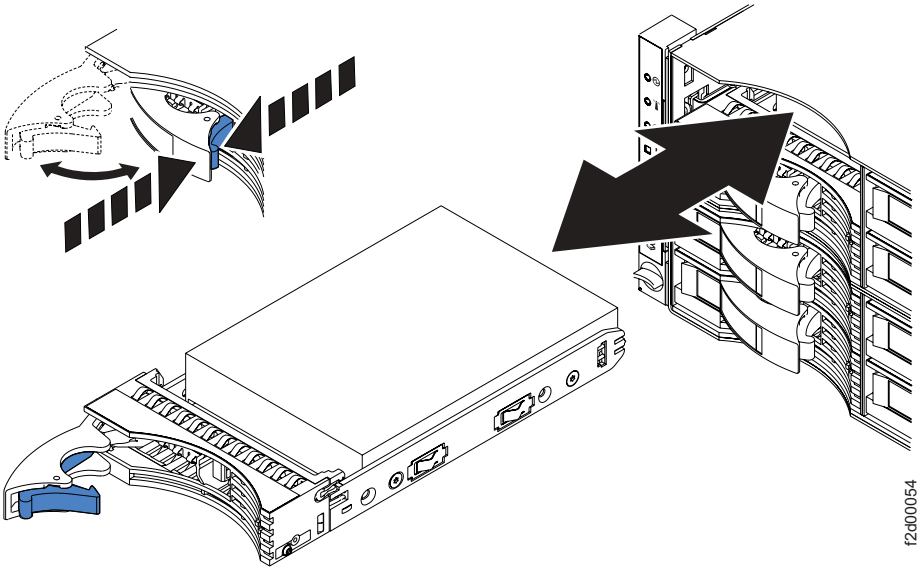

All disk drive module slots should be filled. You can insert a disk drive module blank to prevent overheating the storage unit.

Replacing the power supply

Use approved ESD procedures to prevent damage.

If you removed the power supplies to lessen the weight of the enclosure prior to installation, use the following instructions to replace the power supplies. The power supplies are located on the right and left sides of the rear of the enclosure.

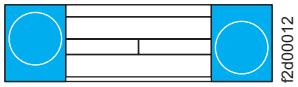

1. Grip the power supply with both hands and align the resource with the slot.
2. Using both hands, push the power supply into the slot until the handle mechanism stops the forward movement.
3. Push the handle in and toward the outer edge of the enclosure until the release button clicks. This inserts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

These images show the steps of a power supply replacement procedure.

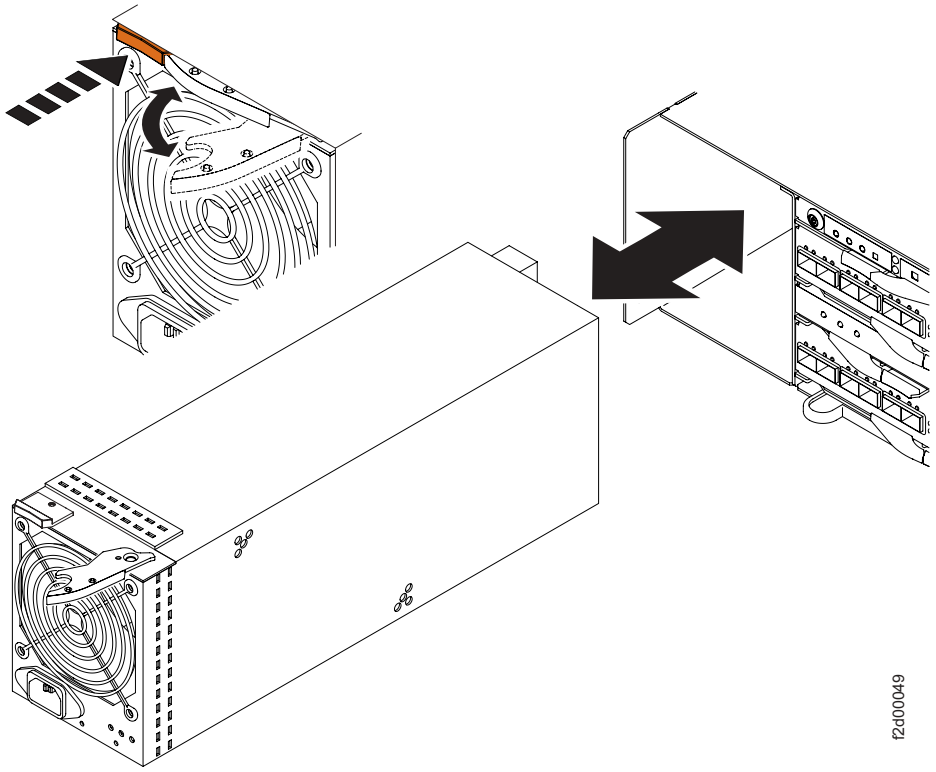

Replacing the rear operator panel

Use approved ESD procedures to prevent damage.

If you removed the rear operator panel to lessen the weight of the enclosure, use the following instructions to replace the rear operator panel. The rear operator panel is located at the top of the rear of the enclosure.

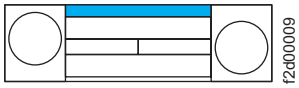

1. Grip the rear operator panel with both hands and align the resource with the slot.
2. Using both hands, push the rear operator panel into the slot until the handle mechanism stops the forward movement.
3. Push the handle in and to the left until the release button clicks. This inserts the remaining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

These images show the steps of a rear operator panel replacement procedure.

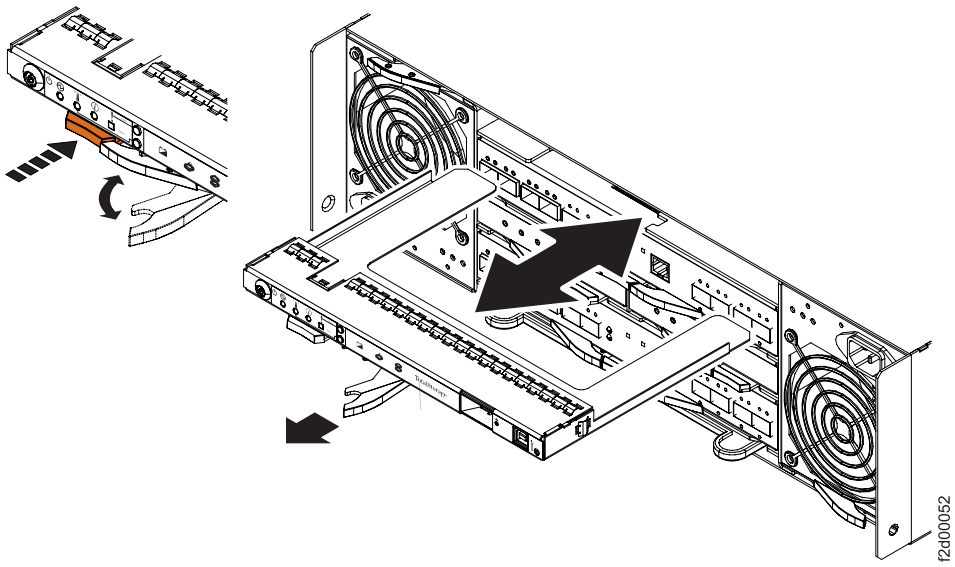

Replacing the battery backup unit

Use approved ESD procedures to prevent damage.

If you removed the battery backup units to lessen the weight of the enclosure, use the following instructions to replace the battery backup units. The battery backup units are located horizontally in the middle of the rear of the enclosure.

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 the steps of a battery backup unit replacement procedure.

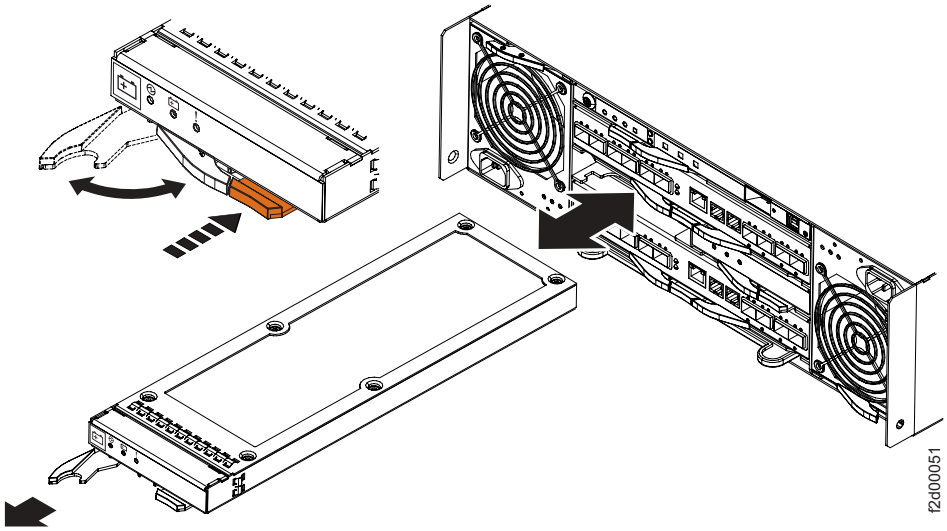

Replacing the processor card

Use approved ESD procedures to prevent damage.

If you removed the processor cards to lessen the weight of the enclosure, use the following instructions to replace the processor cards. 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.

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 the steps of a processor card replacement procedure.

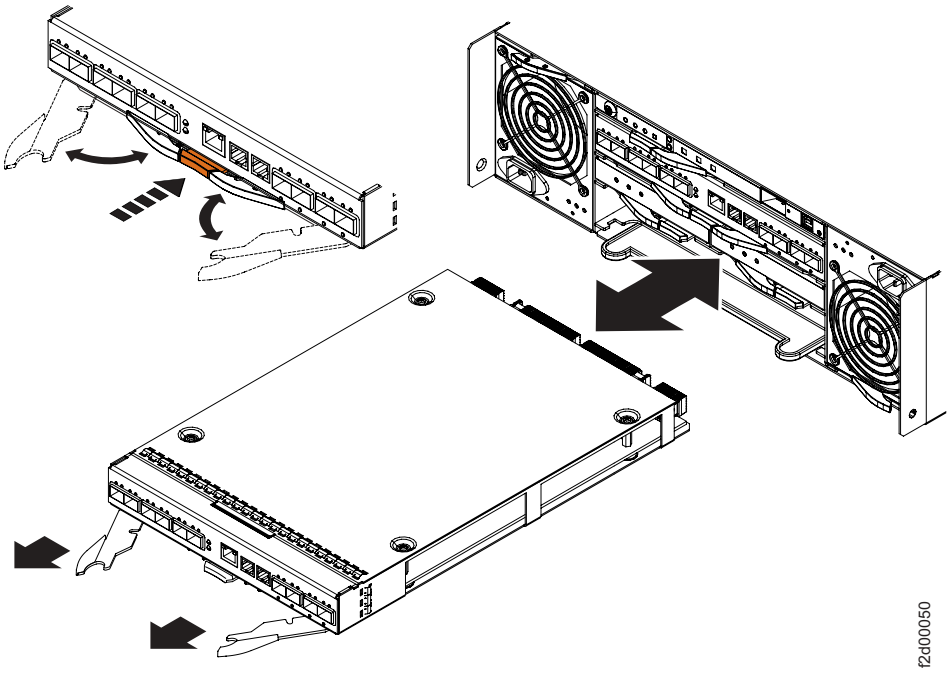

Replacing the front display panel

Follow these steps to install or replace the front display panel or the right front bezel, either when the system is powered off or powered on.

Use approved ESD procedures to prevent damage.

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

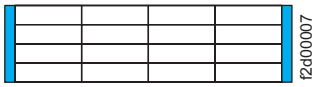

1. Use the handle to align the top of the display panel with the empty slot. The top of the front display panel must fit over the metal support at the top of the slot at the same time that the back of the handle fits under the metal catch on the side of the slot.
2. Push the handle down and into the 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 the steps of a front display panel replacement procedure.

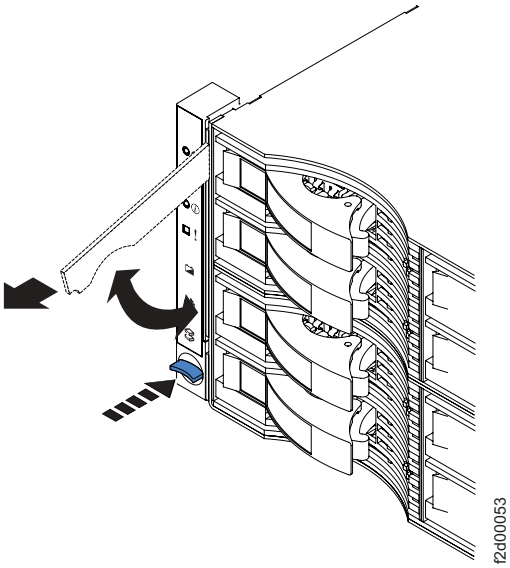

Install the host systems and I/O adapters

1. See the *IBM TotalStorage DS6000 Host Attachment Guide* and the documentation provided with your I/O adapters for installation requirements and procedures.
2. 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-1.ibm.com/servers/storage/disk/ds6000/interop.html*. In addition, for Fibre-Channel host adapters, you can retrieve a list of supported fibre channel HBAs, firmware and device information at the following Web site:*http://knowledge.storage.ibm.com/HBA/HBASearchTool*.
3. Attach fiberoptic 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.

Chapter 3. Routing the cables

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

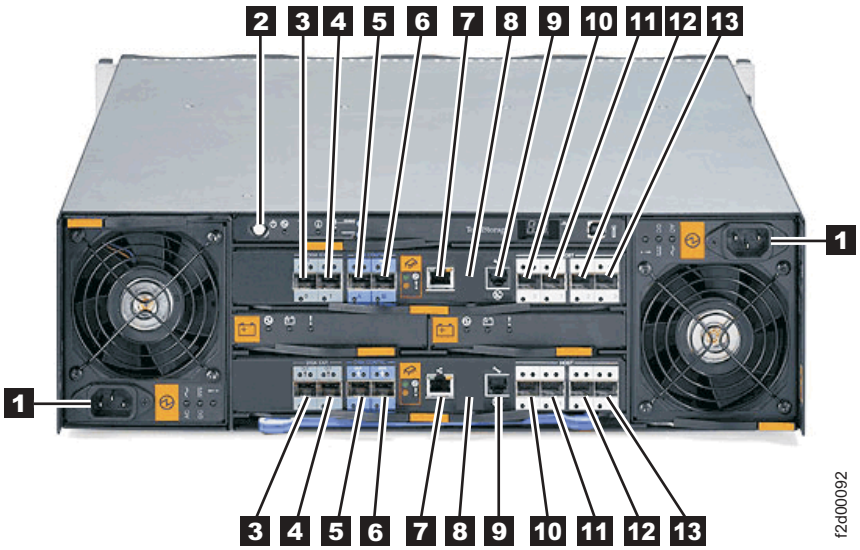

Figure 1. DS6000 server enclosure connection diagram

1. Power supply connector
2. Power control switch
3. 511 processor controller card Disk Exp Port 0
4. 511 processor controller card Disk Exp Port 1
5. 511 processor controller card Disk Ctlr Port 0
6. 511 processor controller card Disk Ctlr Port 1
7. 511 processor controller card Ethernet port
8. 511 processor controller card SCSI Enclosure Services (SES) serial port
9. 511 processor controller card symmetric multiprocessing (SMP) serial port
10. 511 processor controller card Host Port 0
11. 511 processor controller card Host Port 1
12. 511 processor controller card Host Port 2
13. 511 processor controller card Host Port 3

Connect hosts to 511 processor cards

To connect the host adapter to the 511 processor cards, perform the following steps:

1. Install an SFP in a host port on the 511 processor card.
2. Connect the host system cables to the 511 processor card or to switches. The figure below shows the location on the RAID controllers where the host system cables connect. The examples below show

3. Repeat 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.

The figure below shows an example of dual path configuration using fibre-channel switches. Host 1 contains two host bus adapters that are connected to the 511 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 511 processor card and SW2 to the lower 511 processor card.

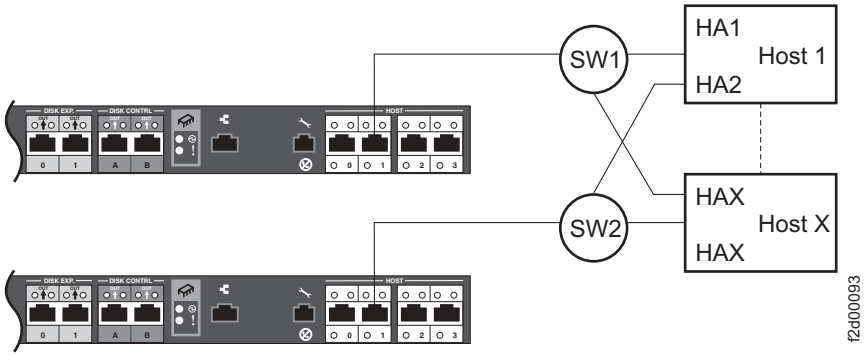

Figure 2. Using two fibre channel switches to connect a host

You can directly attach up to four fibre channel switches to a DS6000. The figure below shows an example of attaching four fibre channel switches.

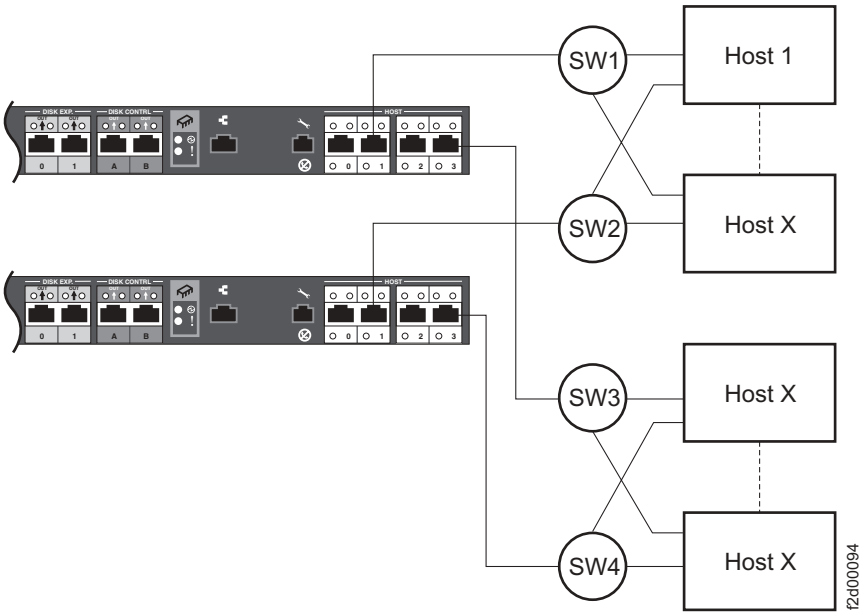

Figure 3. Using four fibre channel switches to connect multiple hosts

Most fibre channel switches support eight to sixteen hosts. If the configuration requires more hosts than four fibre channel switches can support, you must add cascading switches, as shown in the figure below.

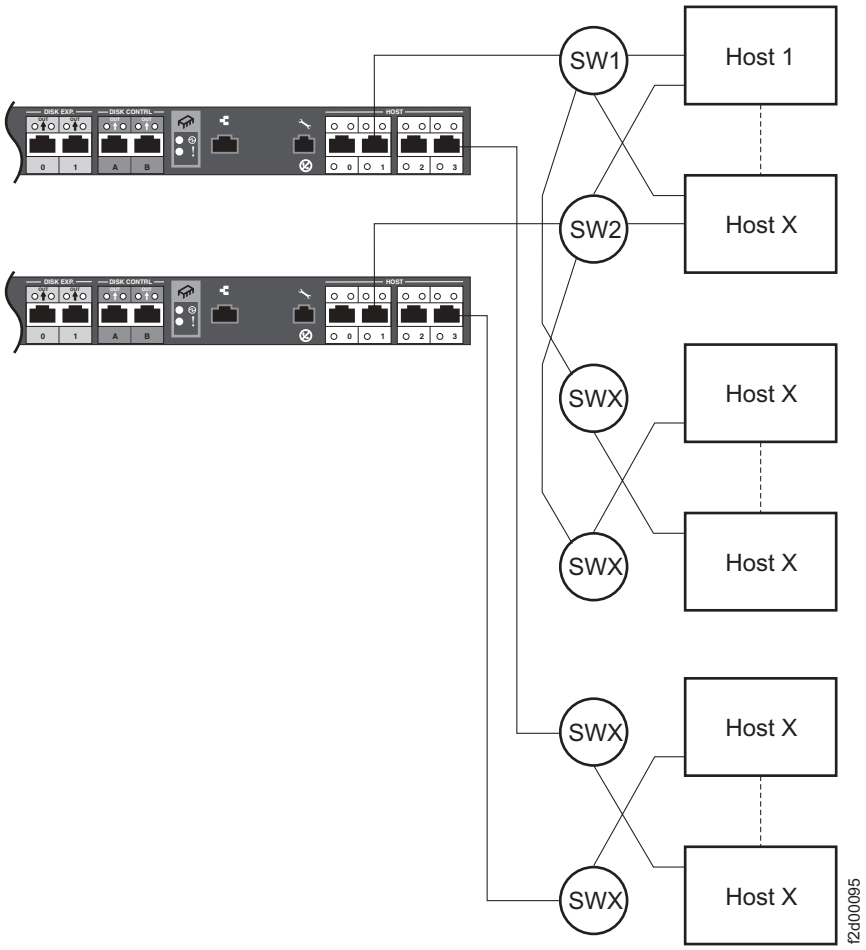

Figure 4. Adding hosts by using cascading switches

Connect secondary interface cables

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.

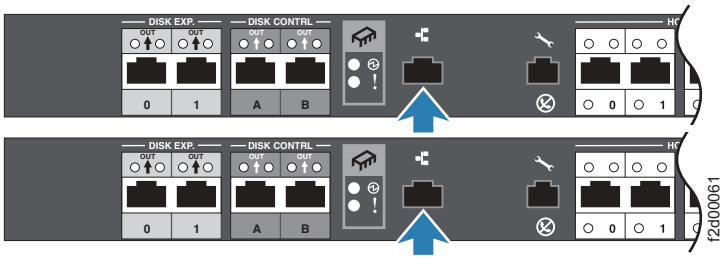

2. The serial interface ports are intended to be used by service technicians to perform diagnostic operations on the server enclosure. Incorrect use of the serial port can result in loss of data access and, in some cases, in loss of data. The processor cards must be able to communicate with each other 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 the Ethernet crossover

cable to connect the two processor cards to each other and the serial cable 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.

Connect storage enclosures

The following are nonconcurrent procedures (the DS6000 is powered down when storage enclosures are connected, as in the case of a new installation).

Note:

1. The recommended sequence for adding up to thirteen storage enclosures is as follows:
   - Connect the first and second enclosures on Loop 1.
   - Connect the third enclosure on Loop 0.
   - Alternate each additional storage enclosure between Loop 1 and Loop 0, starting with Loop 1.

   For example, the first two storage enclosures would be on Loop 1, the third storage enclosure would be on Loop 0, and so on. You can connect up to seven storage enclosures on Loop 1 and up to six on Loop 0 for a total of 13 storage enclosures.

Storage enclosure	Loop
1	1
2	1
3	0
4	1
5	0
6	1
7	0
8	1
9	0
10	1
11	0
12	1
13	0

2. Loop 1 and Loop 0 have different physical cabling connections. Ensure that you follow each procedure exactly.

1. To connect storage enclosures to a server enclosure on Loop 1, perform the following steps:
   a. Use one fiber-optic cable to make a connection from the left `Disk Contrl OUT` port on the upper processor card in the server enclosure to the left `IN` port on the upper processor card of the storage enclosure.
   b. Use one fiber-optic cable to make a connection from the right `Disk Contrl OUT` port on the upper processor card in the server enclosure to the right `IN` port on the lower process card of the storage enclosure.
   c. Use one fiber-optic cable to make a connection from the left `Disk Contrl OUT` port on the lower processor card in the server enclosure to the left `IN` port on the lower processor card of the storage enclosure.

d. Use one fiber-optic cable to make a connection from the right `Disk Contrl OUT` port on the lower processor card in the server enclosure to the right `IN` port on the upper processor card of the storage enclosure.

e. For up to six additional storage enclosures, use two fiber-optic cables to make two connections from the `OUT` ports on the upper storage processor card of the first storage enclosure to the `IN` ports on the upper storage processor card of the second storage enclosure.

f. Use two fiber-optic cables to make two connections from the `OUT` ports on the lower processor card of the first storage enclosure to the `IN` ports on the lower processor card of the second storage enclosure.

g. Repeat steps e through f for up to six storage enclosures.

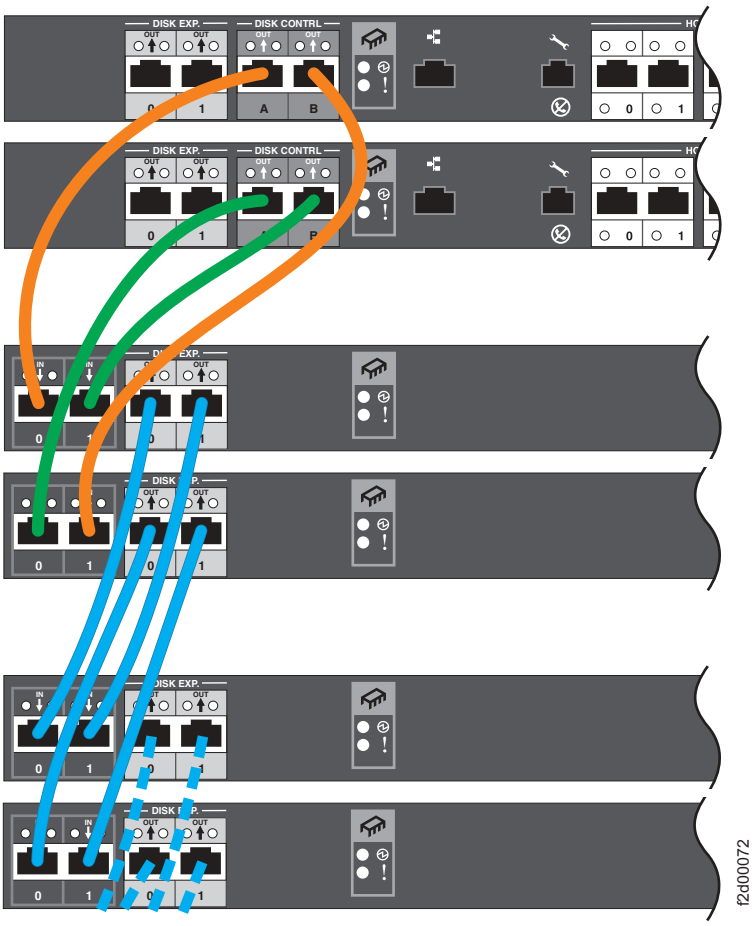

2. Perform the following steps to connect storage enclosures to a server enclosure on Loop 0:

   a. Use two fiber-optic cables to make two connections from the two `Disk Exp OUT` ports on the upper processor card in the server enclosure to the `IN` ports on the upper storage processor card of the storage enclosure.

   b. Use two fibre-optic cables to make two connections from the two `Disk Exp OUT` ports on the lower processor card in the server enclosure to the `IN` ports on the lower storage processor card of the first storage enclosure.

   c. For up to five additional storage enclosures, use two fiber-optic cables to make two connections from the `OUT` ports on the upper storage processor card of the first storage enclosure to the `IN` ports on the upper storage processor card of the second storage enclosure.

d. Use two fiber-optic cables to make two connections from the `OUT` ports on the lower processor card of the first storage enclosure to the `IN` ports on the lower processor card of the second storage enclosure.

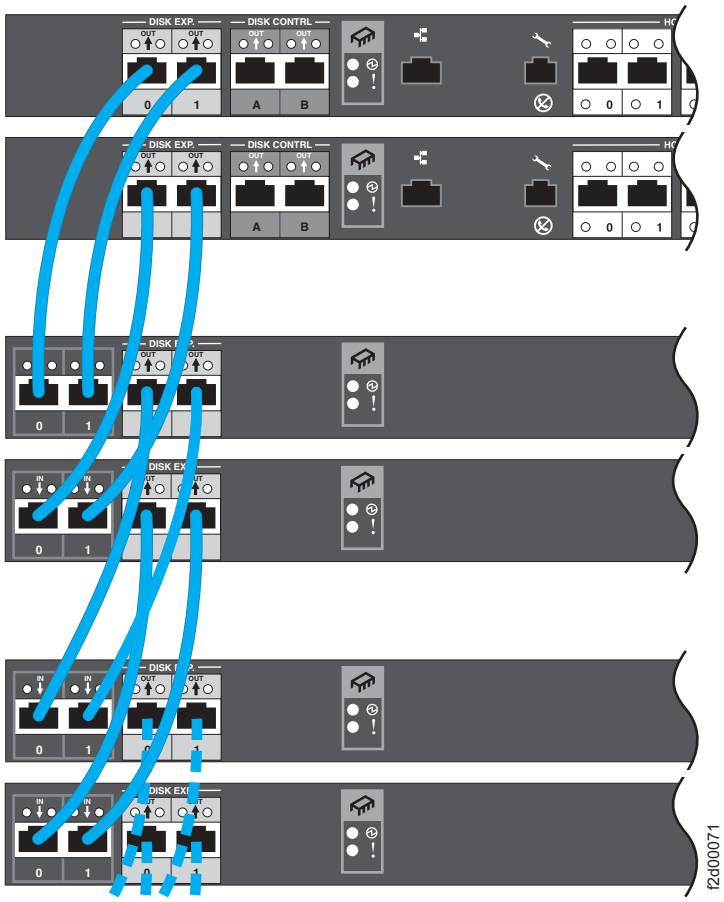

e. Repeat steps c and d for up to five storage enclosures.

Connect power cords

Each server or storage enclosure uses two standard 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.

1. If you have not already done so, attach the power cords to the server and/or storage enclosure(s). The figure below shows the power cord locations.

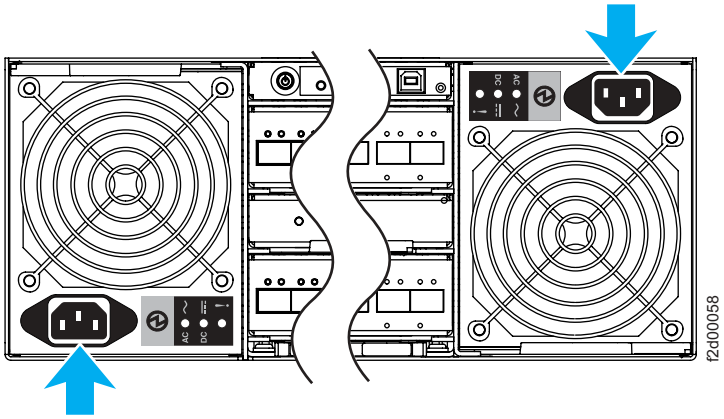

2. Connect the other end of the right server enclosure power cord to a power supply, by plugging it into a properly grounded electrical outlet.
3. Connect the other end of the left server enclosure power cord to a power supply, by plugging it into a properly grounded electrical outlet. To maintain power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

Chapter 4. Checking power and safety

After hardware installation, ensure you can properly power on the server and storage enclosures.

Power on

Note: Prior to power on, ensure the management console has been installed. See Chapter 5, “Installing the DS6000 Storage Manager,” on page 39.

Perform the following steps to turn on the power for the initial startup of the DS6000:

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

   Note: At least four disk drive modules must be present before you start the DS6000.
3. Check the system documentation for the external hardware devices that you intend to power up, and then determine the proper startup sequence. Use the following power-on sequence, where applicable:
   a. Turn on the power to the supporting devices (for example, Ethernet switches and management stations) before the server enclosure.
   b. Turn on the storage enclosures by pressing the Power on/off button (shown below) on each storage enclosure. You must turn on the storage enclosures before the server enclosure. The controllers might not recognize the correct configuration if the storage enclosures are powered up after the server enclosure.

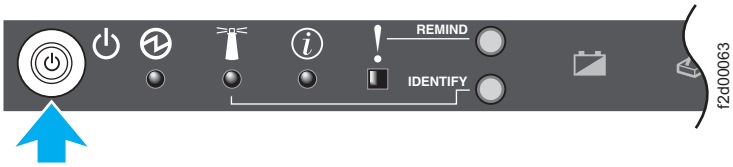

   c. Turn on the power to the server enclosure; then restart or turn on the power to the host.

CAUTION:

Upon power on, a 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.

Interenclosure power operations

You can power the attached storage enclosures on and off through the server enclosure.

Storage enclosure power-on and power-off operations are managed through the server enclosure to which the storage enclosure is attached. You can power on and power off the storage enclosure using the power button on the server enclosure.

Powering on

Initiate the power-on sequence for your storage complex by ensuring that all attached storage enclosures are properly connected. Press the power button that is located on the rear operator panel of the server enclosure. The unit powers on the server enclosure and the attached storage enclosure resources in the necessary sequence. Once the server enclosure has completely powered on, the power-on sequence continues down the line of attached storage enclosures. Each storage enclosure powers on the enclosed resources in the necessary order before the power-on sequence continues to the next attached storage enclosure. Enclosures connected to loop 1 might power on some time after enclosures connected to loop 0.

Powering off

Press and hold for 5 seconds the power button that is located on the rear operator panel of the server enclosure. The unit powers off the server enclosure and each of the attached storage enclosures in sequential order. Once the server enclosure has completely powered off, the power-off sequence continues down the line of attached storage enclosures. Each storage enclosure powers off the enclosed resources in the necessary order before the power-off sequence continues to the next attached storage enclosure. A delay of several minutes can occur before the system fully powers off all the enclosures depending on how many attached storage enclosures are connected.

Verify status through LEDs

The LEDs display the status of the storage 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 checking the LEDs on the front of the storage unit.

1. At the front and rear of the server and storage enclosures, verify the LEDs are displaying the appropriate state.

Table 1. LED status after successful installation

LED Icon	LED Name	Appropriate state
	Power-on	Solid Green
	Location	Solid Blue
	Information	Unlit
	System alert	Unlit

2. If all LEDs are not displaying the appropriate state, refer to the Troubleshooting section to diagnose the problem.
3. You have now completed the hardware installation portion of the DS6000. Continue to the next chapter to complete the installation process.

Setting the IP Address

Use this process to set the IP address

1. Use the serial conversion cable to attach to the serial port on the enclosure. Use a terminal emulator to connect to the server enclosure through the serial port located on the processor card (for example, NetTerm or Windows HyperTerminal). Use the following terminal emulator settings when you connect through the serial port:

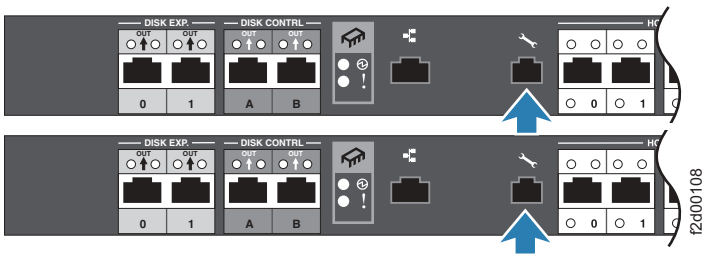

Remote connection setting	Remote connection value
Bits per second	38400
Data bits	8
Parity	None
Stop bits	1
Flow control	Hardware

2. Use the default user ID and password of ″guest″ to access the processor card. The `ncnetconf` script begins automatically.
3. If this is your first time setting the IP address, change the default `guest` password to one that you select.
   a. Choose `Change 'guest' password` from the ncnetconf 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.
4. Choose `Configure network parameters` from the ncnetconf Main Menu options.
5. To set the IP address, perform the following steps:
   a. Choose `Use static IP address` from the Network configuration menu options.
   b. To change the IP address for this node, choose `IP address for this node` from the Static IP addresses configuration menu options. To change the IP address for the other node, choose `IP address for other node` from the Static IP addresses configuration menu options.

      Note: The ncnetconf program prevents you from setting a network mask IP address that might conflict with the above addresses.
   c. When the `IP Address?` prompt appears, type in the appropriate IP address and press Enter.
   d. Select `Back to Network Configuration` to return to the Network configuration menu.
   e. Select `Advanced Configuration Options` to set the domain name server and the gateway settings.

f. Select `Back to Network Configuration`

g. Select `Back to Main Menu` to return to the ncnetconf Main Menu.

h. Select `Apply changes and exit` from the options in the main menu to save your changes and exit the application.

Note: Ensure that you set the IP address and that it is cabled to the same subnet as the DS6000.

6. Repeat Step 4 for the second node.

Multiple IP addresses on the DS6000 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.

Use the following procedure to change the binding order for the IP address:

1. From the **Control Panel**, open the **Network Connections**.

f2d00121

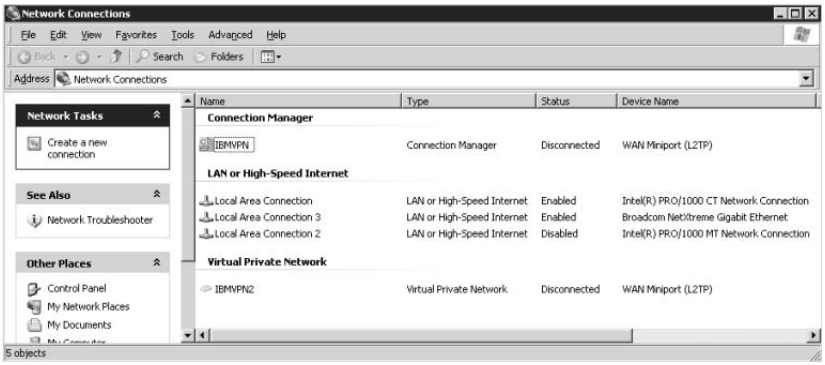

f2d00122

2. From the **Network Connections**, select **Advanced** from the menu; then select **Advanced Settings**.

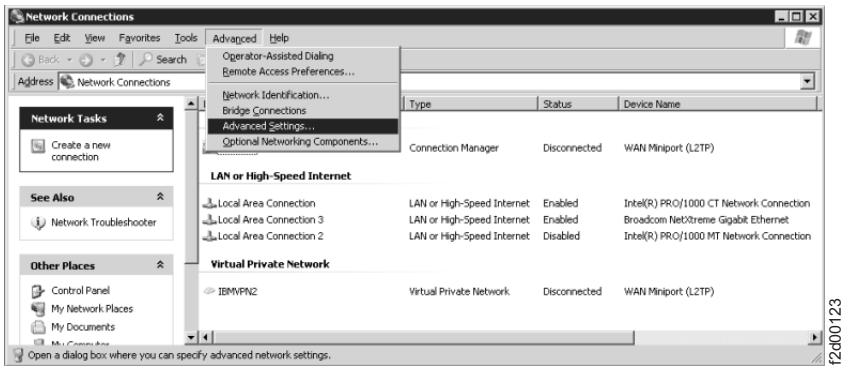

3. In the **Adapters and Bindings** tab, check the list of the network adapters to ensure that the first network adapter is the adapter that is on the same subnet network as the DS6000. If this is not the case, you must change the binding order so that the IP address on the same subnet (private) as the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

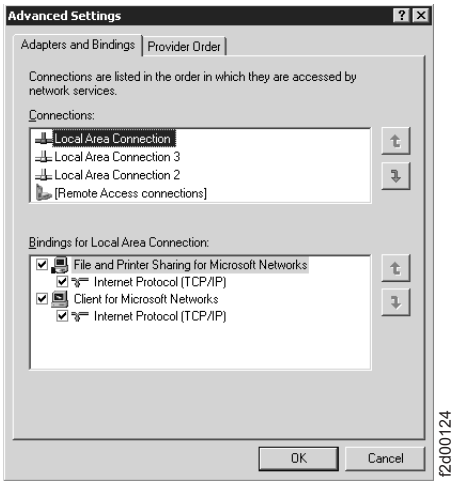

Chapter 5. Installing the DS6000 Storage Manager

The DS Storage Manager is installed through a GUI or in silent mode for Windows operating systems. It can be accessed 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 receive access to the use of copy services functions if you have the authorized license management code.

You can also install the DS6000 Storage Manager on a PC that has the following minimum footprint requirements. Also, ensure you have installed a supported browser.

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

Supported browsers:

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

Supported operating systems:

Operating System	Full management console install	Offline management console install
Windows Server 2003 Enterprise Edition	X	X
Windows Server 2003 Standard Edition	X	X
Windows 2000 Advanced Server SP4	X (English only)	X
Windows 2000 Server SP4	X (English only)	X
Windows 2000 Professional SP4	X (English only)	X
Windows XP Professional SP1		X
Windows XP Professional SP1a		X
Windows XP Professional SP2	X	X

Note: The DS6000 Storage Manager is not supported on any Windows 64-bit operating system.

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

- Internet Explorer
  1. From the **Tools** menu select **Internet Options**.
  2. Select the **Advanced** tab and scroll down to the **Multimedia** section.
  3. Check **Play animation in web pages**.
- Netscape

1. From the **Edit** menu select **Preferences**.
2. Double-click on **Privacy and Security**.
3. Select **Images** and select **as many times as the image specifies** in the **Animated image should loop** section.

Ensure that you have available these completed planning work sheets, which can be found in the *DS6000 Introduction and Planning Guide*, Chapter 5 ″Planning your DS6000 series″:

- Configuration work sheet
- Network settings work sheet

Ensure that your Windows PC has the regional settings option specified to your language and Country or Region.

Terminal Server:

Use these procedures if your Windows PC has the Terminal Server enabled and the session that you are logged in 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.

Installing the IBM TotalStorage DS Storage Manager on the Windows operating system

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

You can choose to install the IBM TotalStorage 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 mode (also called silent mode) – allows you to customize a response file and issue a command to complete the installation process.

After you have installed the DS6000 Storage Manager, the following results occur:

- Activation of the IBM TotalStorage DS Storage Manager server and the IBM TotalStorage DS Network Interface server. These servers are set to automatic startup so that when you start your computer these servers are automatically activated.

- Activation of the DS Storage Manager application, which includes the real-time and simulated manager components. These components can both be installed on the same machine and are integrated into the user interface. They are designed to help you create and manage the physical and logical configurations of your storage complexes 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 IBM TotalStorage DS Storage Manager on the Windows operating system using the graphical mode

Use the following steps to install the IBM TotalStorage DS Storage Manager in your Windows environment using the graphical mode. Installation can take approximately from 10 to 20 minutes.

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

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

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

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

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

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

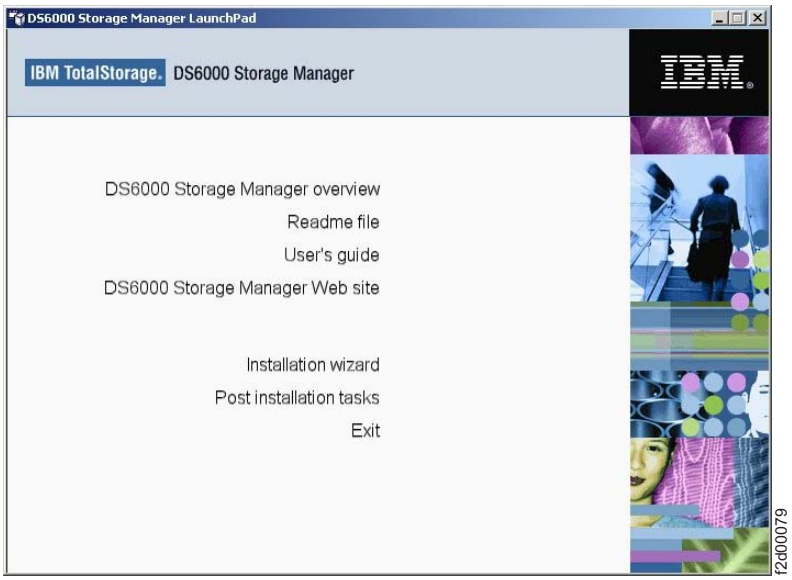

Launchpad window

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

   DS6000 Storage Manager overview
   Provides information about the IBM TotalStorage DS Storage Manager software.

   Readme file (recommended selection)
   Provides last minute product information that was not provided in these installation instructions.

   User's guide
   Provides specific installation instructions.

   DS6000 Storage Manager Web site
   Provides information from the product Web site.

   Installation wizard
   Starts the IBM TotalStorage DS Storage Manager installation program.

   Post installation tasks
   Provides information about configuring the IBM TotalStorage DS Storage Manager.

   Exit
   Exits the IBM TotalStorage DS Storage Manager LaunchPad program.

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

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

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

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

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

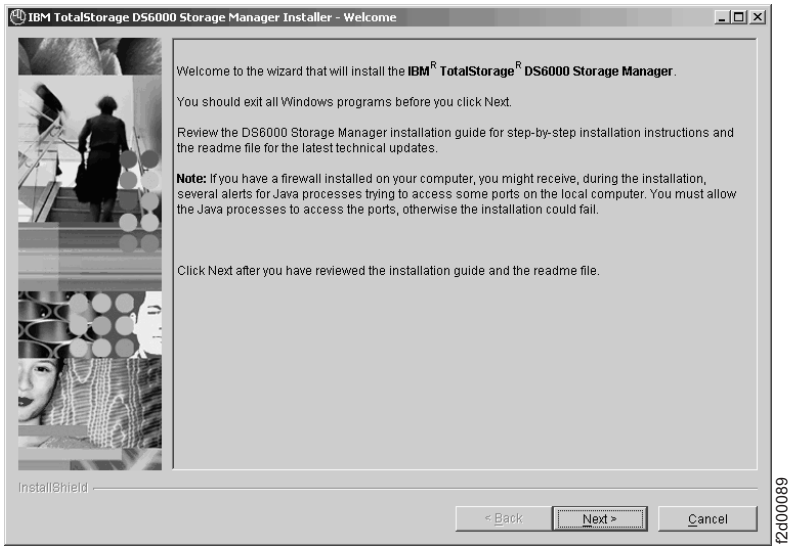

DS6000 Storage Manager Installer Welcome Window

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

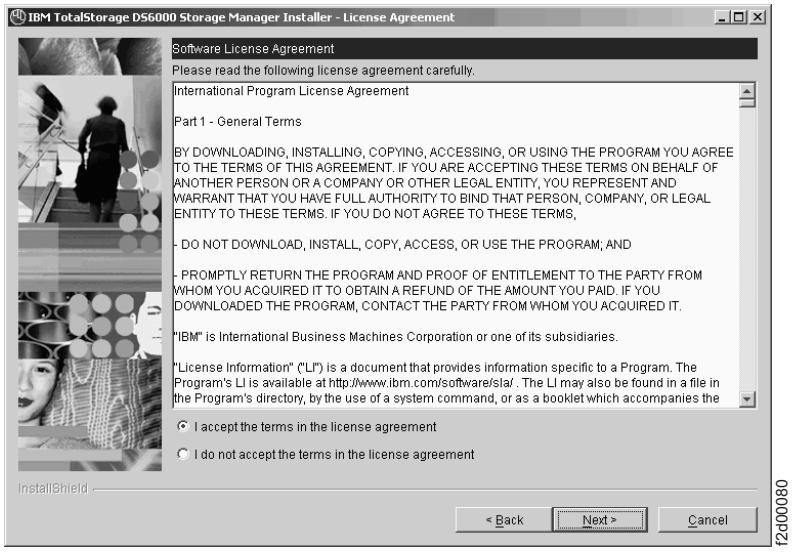

License Agreement window

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

7. Read the license agreement and click your acceptance or nonacceptance of the agreement. If you accept, the **Next** button is highlighted. Click **Next** to

continue or click **Cancel** to exit the installation. When you click **Next**, the Destination Directory window is displayed.

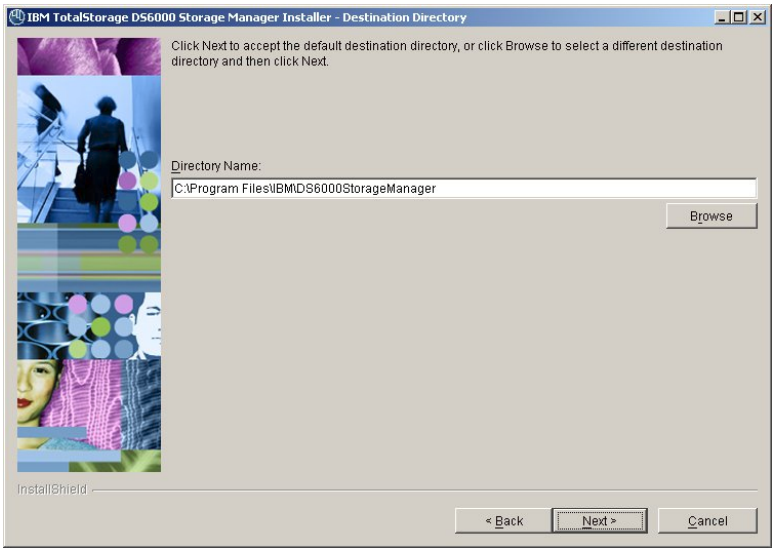

Destination Directory window

8. Choose the directory on the Destination Directory window where you want to install the application by using one of the following options. Clicking **Next** after any one of the options causes the Management Types Selection window to display.
   a. Click the **Next** button to accept the default (recommended) directory that is shown in the window.
   b. Type a fully qualified directory path into the Directory Name path to designate your own directory. Then click the **Next** button.
   c. Click the **Browse** button and a directory window is displayed. Scroll to the directory that you want to use for the application. Click it and click **OK**. The directory path is displayed in the Directory Name field. Click **Next** to continue.
   d. Click **Cancel** to exit the installation process.
9. Select the management type that you want to install by selecting one of the following options on the Management Type Selection window:
   a. **Full Management Console install**

      Note:
      - You cannot install this management type on a storage unit that already has an existing DS Storage Manager.
      - You cannot install this management type on a non-English version of the Windows 2000 operating system. (The VPN used by this installation does not work on a non-English version of the Windows 2000 operating system.)

   b. **Offline Management Console install**

      Note: This management type can be installed on a storage unit that already has an existing DS Storage Manager.

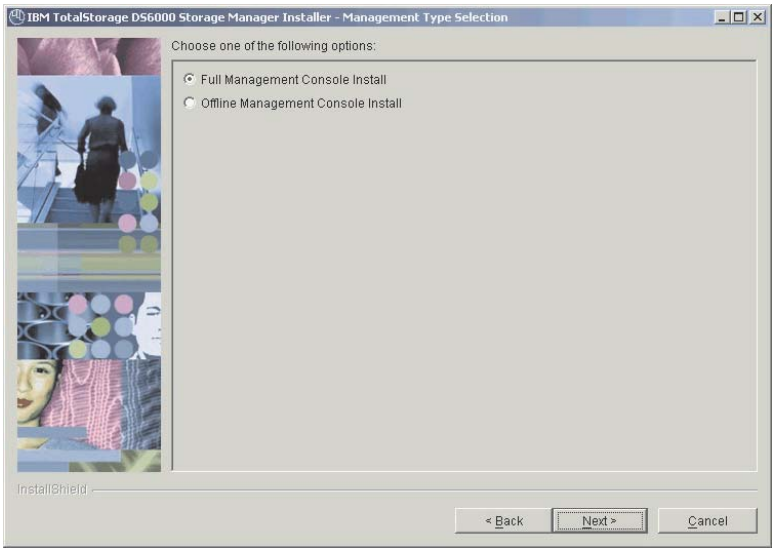

Windows Management Type Selection window

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

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

      When you click **Next**, the Installation Confirmation window is displayed. If your installation process has taken you through these two checking windows, go to Step 16 on page 48.
    - The Server Parameters window is displayed with default values shown for all the fields.

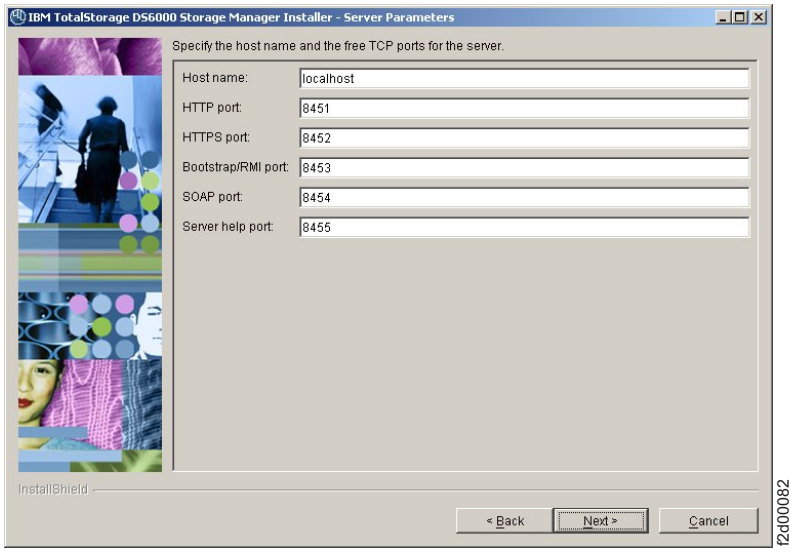

Windows Server Parameters window

11. Click **Next** after you have ensured that all the input fields on the Server Parameters window are complete. The SSL Configuration window is displayed. This window allows you to supply the fully qualified name of the two server key files generated before or during the DS6000 Storage Manager installation.

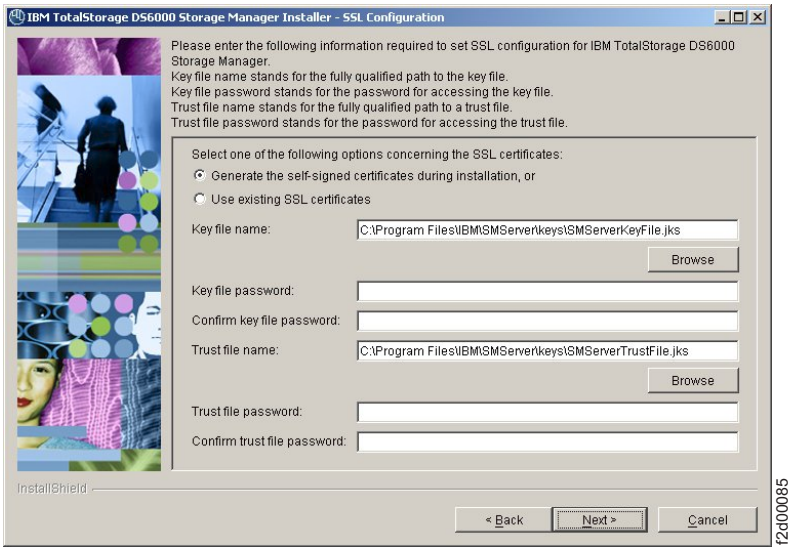

Windows SSL Configuration window

12. Follow the instructions on the window.
    a. First, designate whether you want to generate a new SSL certificate by selecting **Generate the self-signed certificates during installation**. This is the most likely selection when you are installing DS6000 Storage Manager for the first time. You can also select **Use existing SSL certificates**. This choice is most likely made when the key files already exist and you do not want to generate new key files.
    b. Complete the rest of the fields on the window. Provide the fully qualified path to the key file and trust file by using the Browse button. The key file and trust file are identified by an extension of .jks.
    c. Supply a password for each file. The password must be a minimum of six-characters and it cannot contain any blanks.

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

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

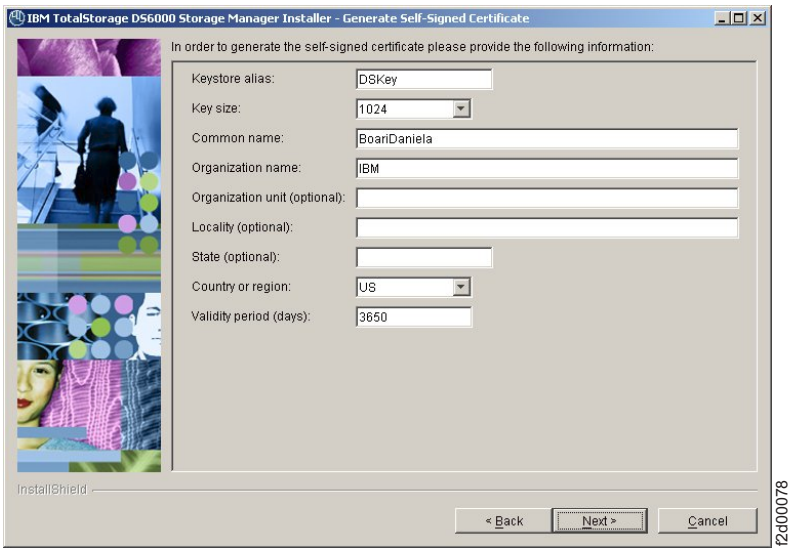

Windows Generate Self-Signed Certificate window

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

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

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

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

    Organization name
    The name of your company

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

    Locality (optional)
    The city or location where your company resides.

State (optional)
The state or province where your company resides.

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

Validity period (days)
Specifies the lifetime of the certificate in days.

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

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

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

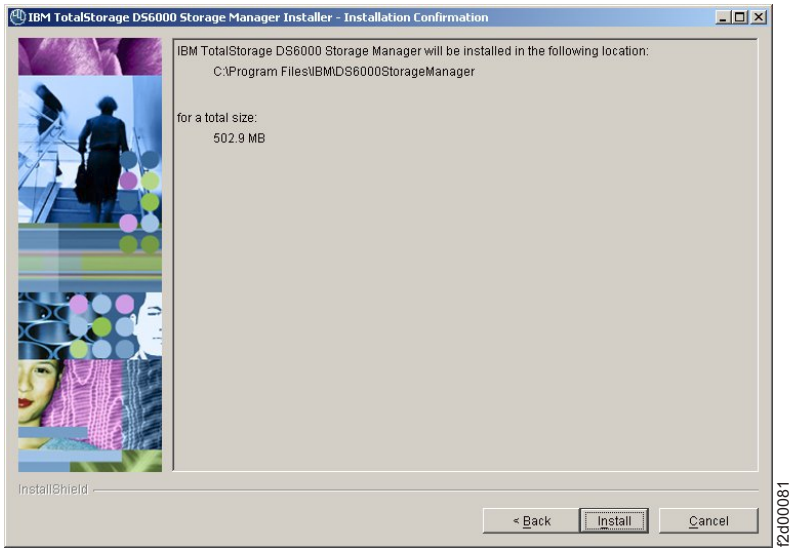

Windows Installation Confirmation window

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

    The installation process performs the following actions:

    a. If the two servers (DS Storage Manager Server and DS Network Interface Server) are already installed on your system, they are stopped in the following order of windows:

        1) The Embedded IBM WebSphere Application Server - Express server (part of the DS Storage Manager Server) window is stopped first.
        2) The service window (DS Network Interface Server) is stopped next.
        3) The WS Help System (part of the DS Storage Manager Server) window is stopped, if it was not stopped before by the Embedded IBM WebSphere Application Server.

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

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

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

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

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

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

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

17. Click **Finish** to exit the installation process. When the installation process exits, a file (postinstallation.txt) is displayed, containing the post installation tasks. The post installation and configuration tasks are also in the *IBM TotalStorage DS6000 Installation, Troubleshooting, and Recovery Guide*. Use these instructions to complete the setup tasks.

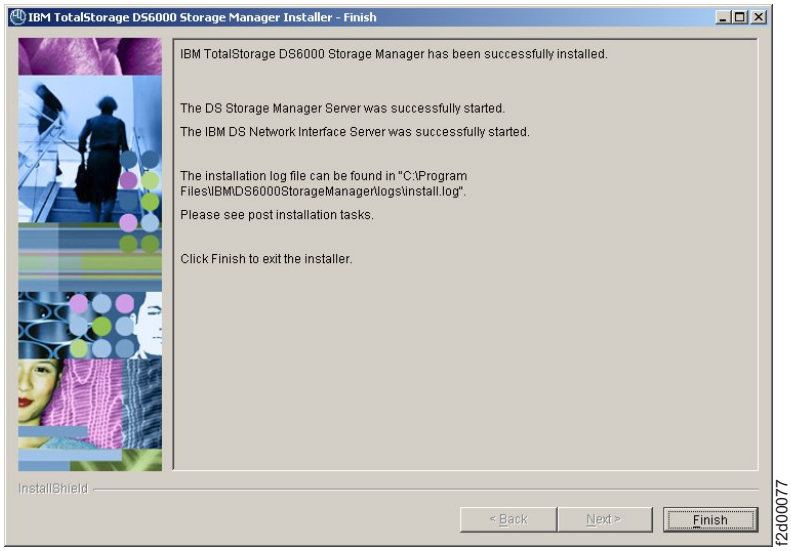

f2d00077

Windows DS6000 Storage Manager Installer Finish window

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

18. Restart to complete installation.
19. If the Finish window indicates that the installation fails, check the installation log for error messages. The installation log is located in *xxx\logs\install.log*, where *xxx* is the destination directory where the IBM TotalStorage DS Storage Manager is installed (for example, c:\Program Files\IBM\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 silent mode of installation.
    - If the error is due to a wrong installation parameter value, restart the installation using the steps that are described in this procedure or the steps that are described in the silent mode of installation. Navigate to the Server Parameters window and insert the correct values. Then finish the installation process.
20. Complete the postinstallation tasks. If, when exiting the installation wizard, the postinstallation.txt file does not automatically open, manually open it from the LaunchPad window, and follow the instructions to complete the postinstallation tasks.
21. Exit the LaunchPad program by clicking **Exit** on the LaunchPad window. Reboot if instructed to do so.
22. You have now completed the software installation portion of the DS6000. Continue to the postinstallation tasks.

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

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

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

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

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

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

1. Log on to your Windows system as an administrator.

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

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

   a. Ensure that your command prompt is located at the root directory of the CD-ROM drive.

   b. Type `setup.exe -options-template <responsefile-path>\<responsefile>`, where *responsefile-path* is the path where you want the response file to be created, and *responsefile* is the name of the response file that you want to create.

   c. Press the **Enter** key and wait for the program to create the response file.
4. Open the response file using a text editor and modify the default options with the values that you want.

   Follow these guidelines when you modify the default options:

   - Remove the three # characters from the beginning of each line that contains a parameter (if you generated the response file as shown before). Change the parameter default value to the value that you want for that option. You *must* enclose all values in double quotation marks (″<value>″).
   - Verify that all paths generated or modified in the response file are correct. There are some cases when the generated file displays an incorrect string representing the path of the installation location. For example, the parameter ″-P installLocation = ″C:Program FilesMy Product″″ is wrong and must be ″-P product.installLocation = ″*<directory where you want the DS Storage Manager installed>*″″
   - Ensure that the entered values for the ports values are available and not used by other applications on your system. Use the command ″*netstat -a*″ to verify the ports in use on your system.
   - Specify two passwords for the SSL keys files. These passwords must observe the following criteria:
     - The passwords must contain at least 6 characters.
     - Use the following options within the responsefile.txt file to set these options:
       - For the keystore use:

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

       - For the truststore use:

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

         Note: These options do not have default values. If you do not set these values, the installation process fails.
5. Save the modifications to the response file.
6. Type the following command at the command prompt and press the **Enter** key on your keyboard to start the installation process in silent mode: *setup.exe -options <responsefile-path>\<responsefile> -silent* where
   - *responsefile-path* — represents the path where the response file resides.

- *responsefile* — represents the relative name of the response file that you used or created (for example, responsefile.rsp or responsefile.txt).

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

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

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

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

```
##############################################################################
#
# InstallShield Options File 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 TotalStorage 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"
```

Chapter 6. DS Storage Manager postinstallation instructions

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

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

Internet browser support

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

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

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

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

Starting the DS Storage Manager program

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

The DS Storage Manager program times out over a period of inactivity or when you turn off your computer.

Starting the DS6000 Storage Manager on a Windows operating system

You can open the DS6000 Storage Manager using the Programs list. The DS6000 Storage Manager opens in your default browser.

Log on to your Windows operating system and use the following steps to access the DS6000 Storage Manager program.

1. Click **Start**.
2. Click **Programs**.
3. Click **IBM TotalStorage Manager 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. You will be prompted to install the security SSL certificate. Click **Yes**.

When you start the DS6000 Storage Manager, the IBM TotalStorage DS6000 Signon window is displayed. Enter the user name and password to access the program. The default user name is **admin**, and the default password is **admin**. The first time you log on using the default user name and password, a second login screen appears. Change your password on a the second login screen.

Stop and start the DS Storage Manager Servers

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

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

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

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

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

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

Getting started with the DS Storage Manager console

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

The following components are available to you within the DS6000 Storage Manager program:

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

All of the functions that are associated with these components are not described here. If you need help with any of the pages that you open, you can click the question mark icon on the tool bar. Page help is displayed in another window along with access to an Information Center that contains information about the DS6000.

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

Log In

You can log in to the DS6000 Storage Manager by providing your user name and password. This function is generally administered through your system administrator and by your company policies.

Note:

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

Simulated component

If you have installed the DS6000 Storage Manager in Full-management console or offline mode, you can begin using the simulated component immediately after logging into the DS Storage Manager. This component provides the ability to create or modify logical configurations when your storage unit is disconnected from the network. After creating the configuration, you can save it and then apply it to a network-attached DS6000 at a later time.

Real-time component

You can use the Real-time component selections of the DS6000 Storage Manager if you chose **Full Management Console install** during the installation of the DS6000 Storage Manager. This application provides the ability to work with the logical configuration and Copy Services features when attached to a DS6000 network.

Copy Services

You can use the Copy Services selections of the DS6000 Storage Manager if you chose the **Full Management Console install** option during the installation of the DS6000 Storage Manager and you have purchased these optional features.

Chapter 7. Creating a configuration

This section contains information for creating either a real-time or simulated configuration.

DS6000 Configuration overview

The storage that is associated with a logical device is called a logical volume.

A logical device with nonremovable media is associated with only one logical volume. A logical volume consists of one or more extents, where each extent is identified with a contiguous range of addressable data units on the logical volume.

For internally configured storage, a group of storage devices are associated with an array site. One or more array sites creates an array. A combination of one or more arrays create a logically contiguous storage space called a rank. On a multiarray rank, striped data (RAID 0) improves performance across the arrays in the rank.

For externally configured storage, one or more storage devices in an external storage node creates one or more external LUNs. Each external LUN is a logically contiguous storage space that is also called a rank. Each rank is divided into fixed sized extents. A set of extents made up of one or more ranks is called an extent pool. One or more extents from the same extent pool are allocated to a logical volume to create the logical storage space that is associated with a logical device.

In most cases, extents are assigned to a logical volume when the volume is created. In certain cases, you can dynamically allocate extents to the logical volume because the extents store data that is written to the logical volume.

Creating real-time configurations

The topics in this section provide information for creating real-time storage configurations.

Creating an express real-time storage configuration

Use this process to create an express real-time storage configuration.

You must have your storage unit installed. See “Creating a storage unit (Simulated only)” on page 65 for more information.

This scenario covers one method for creating a configuration on the storage unit. Your network must be configured so that it can support all the components and functions that you will use with your storage unit.

1. Assign the storage unit to a storage complex. In the navigation, under **Manage Hardware**, select **Storage Complexes**. Select a storage complex from the table. See “Assigning a storage unit to a storage complex” on page 66 for more information.

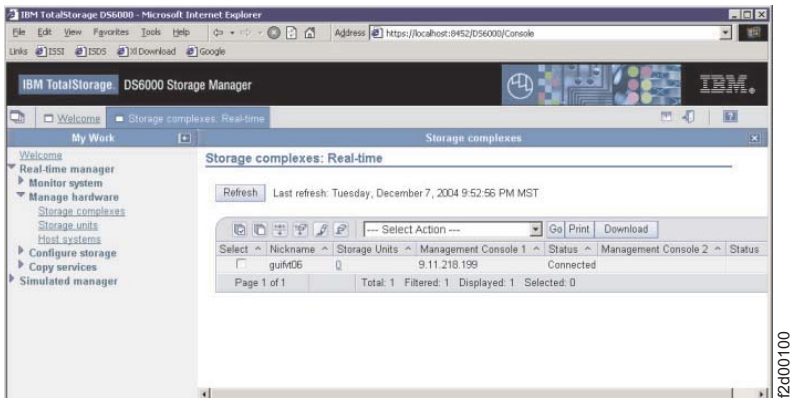

Note: At this point, you can now install and use the DSCLI to complete the remaining post installation tasks, or you can continue using the DS Storage manager. See “Using the DS CLI application” on page 68 for more information.

2. Activate licenses. In the navigation, under **Real-time manager**, select **Manage Hardware**, then select **Storage Units**. See “Activating licenses” on page 71 for more information.
3. (Optional) Define a peer storage complex by clicking on **Storage complexes** and selecting a storage complex from the table. See “Defining Multiple Management Consoles” on page 73 for more information.
   a. Click **Define Peer** in the **Select Action** drop-down list to assign the storage unit to the selected storage complex, and then click **Go**.

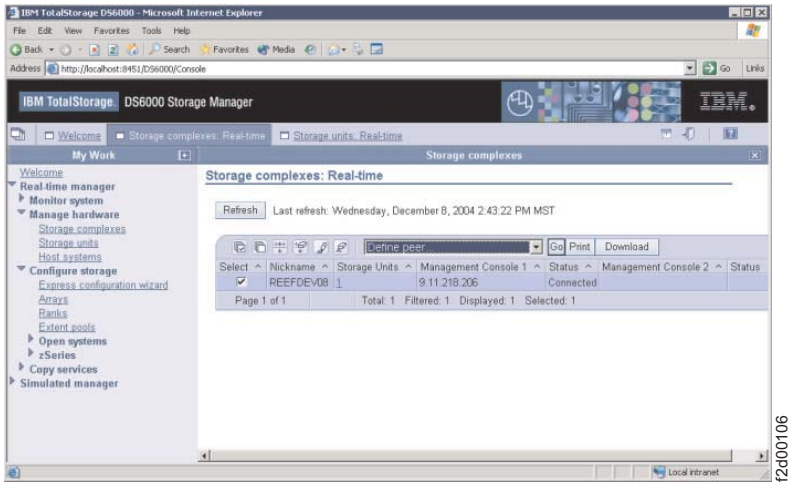

   b. Define either the hostname or IP address (not both) for the peer management console. Click **OK** when you are finished.

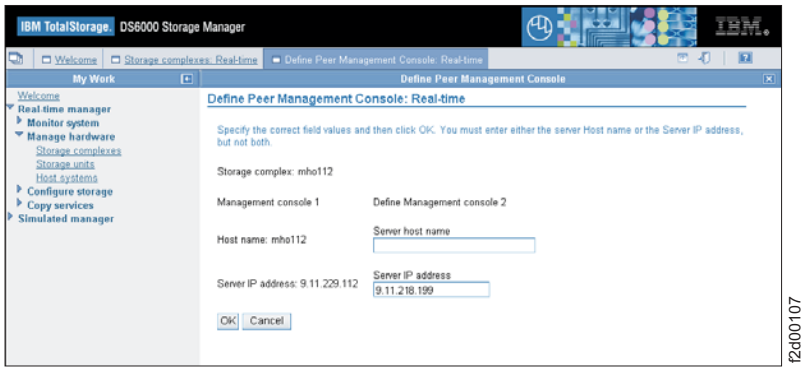

c. Click on the **Storage Complexes** main page and you should see a second management console nickname defined with status on the right most columns. See “Assigning a storage unit to a storage complex” on page 66 for more information.

4. Configure the storage unit. In the navigation, select **Manage Hardware**, then **Storage Units**. Select a storage unit in the table. In the **Select Action** drop-down box, select **Configure** and then **Go**. See “Specifying storage unit day and time” on page 74 and “Reviewing storage unit network settings” on page 74 for more information.
5. Define customer contact information. In the navigation, select **Manage Hardware**, then **Storage Units**. In the **Select Action** drop-down box, select **Customer contact** and then **Go**. See “Defining customer contacts” on page 74 for more information.
6. Configure notifications. In the navigation, select **Manage Hardware**, then **Storage Units**. In the **Select Action** drop-down box, select **Configure notifications** and then **Go**. See “Setting up call home” on page 123 for more information.
7. Configure I/O ports for the host system. In the navigation, select **Real-time Manager** → **Manage hardware** → **Storage units**. Select the storage unit to configure. In the Select Action drop-down box, select **Configure I/O Ports...** and then **Go**. See “Configuring I/O ports” on page 74 for more information.
8. Create the host system. In the navigation, select **Manage Hardware**, then **Host Systems**. In the **Select Action** drop-down box, select **Create** and then **Go**. See “Creating host systems” on page 75 for more information.
9. Configure storage. In the navigation, select **Real-time manager**, **Configure Storage**, then **Express Configuration Wizard**. See “Creating open systems volumes” on page 76, “Creating iSeries volumes” on page 77, and “Creating zSeries volumes” on page 78 for more information.

   Note: The express configuration tool defines and configures for only one host, not multiple hosts.

Creating a storage unit (Simulated only)

Use this process to create a storage unit and to specify its attributes and properties.

1. In the navigation, select, in order, Simulated Manager, Manage Hardware, and Storage Units. In Storage Unit — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create Storage Unit — General storage unit information page is displayed.
2. 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 need 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.

3. Click **Next** to continue.
4. The Create Storage Unit —Specify DDM packs page is displayed. You must specify the **Quantity of DDM packs** and the **DDM type**. Click **Add**, then click **Next** to continue. The Define licensed function page is displayed.
5. In the Define licensed function page, you must specify a value in the **Operating Environment License (TB)** field.
   a. 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.
6. Specify values in the remaining four fields as appropriate. Click **Next** to continue.
7. The Verification page is displayed. Use this page to review the established 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 storage unit creation process.

Assigning a storage unit to a storage complex

Use this process 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.

1. In the navigation, select, in order, Real-time Manager, Manage Hardware, and Storage Complexes. In Storage Complexes— Main Page, select the appropriate storage complex from the table.

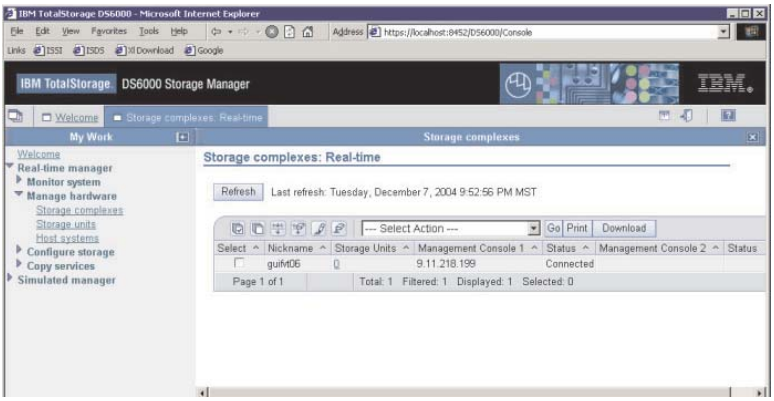

f2d00100

2. In the **Select Action** drop-down list, select **Assign Storage unit** and then **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 serial number. The serial number must be at least 7 digits, with the first 2 digits being the point of manufacturer (for example, if the serial number is 13AAXRA, the point of manufacturer is 13).

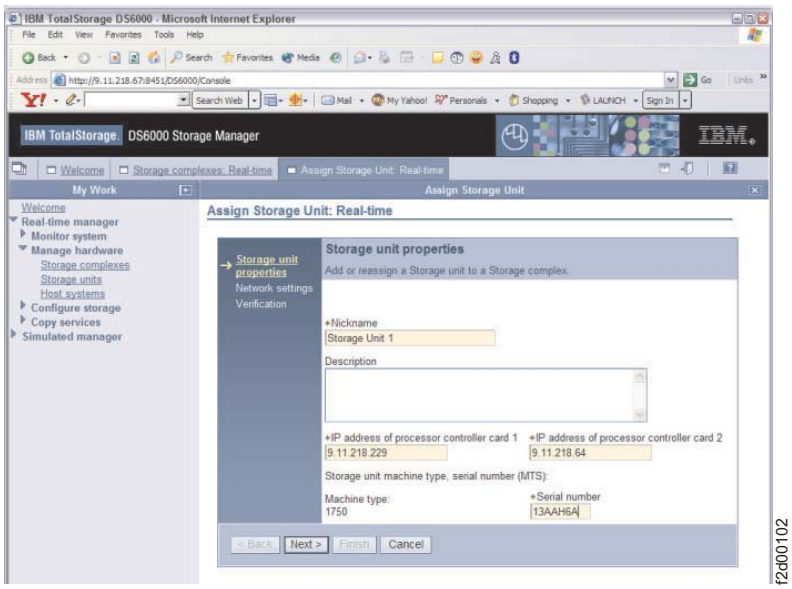

3. Click **Next**. The Network settings page is displayed.

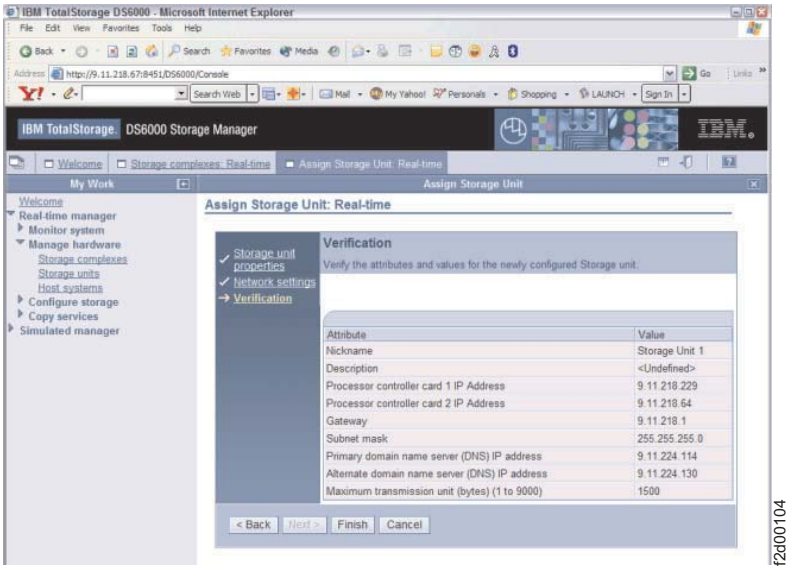

4. Specify the appropriate network settings and then click **Next**. The Verification page is displayed.
   a. Enter a gateway
   b. Enter a subnet mask
   c. Enter the primary DNS address
   d. Enter the secondary DNS address
   e. Enter a different Max transmission units value, if necessary.

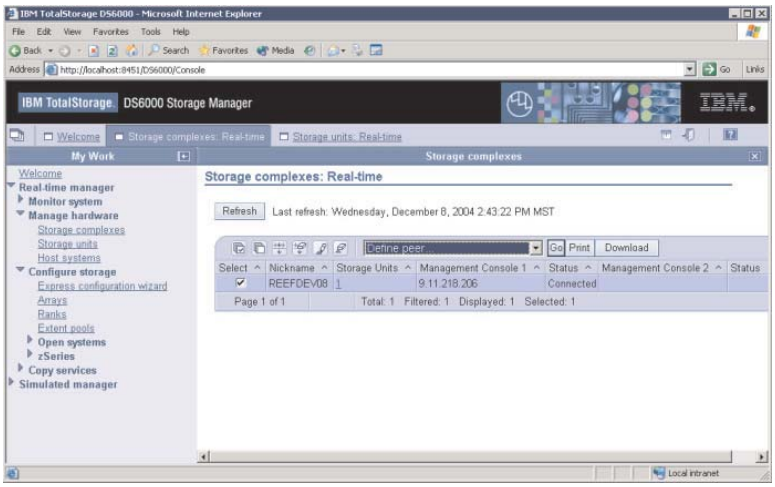

f2d00106

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

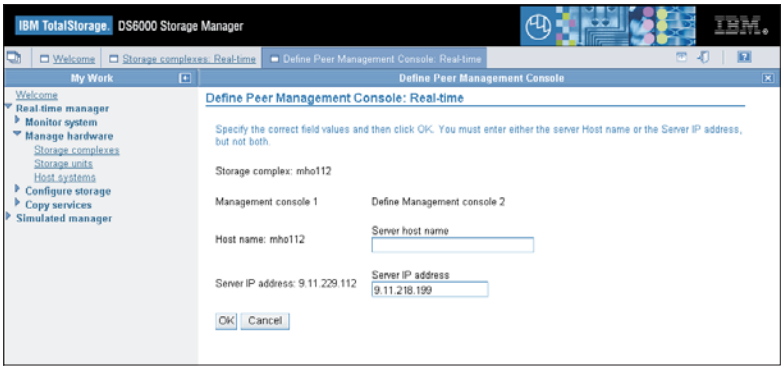

f2d00107

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, there are three command modes that are available to you.

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

- Single-shot
- Interactive
- Script

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

Sun Solaris
/opt/ibm/dscli

Windows
C:\Program Files\IBM\dscli

HP Tru64
/opt/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 is expired when 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`.
- Supply the log in information in a profile configuration file (for additional information, see the topic ″Default configuration setup with a profile file″). 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.

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. Use the following example to use the single-shot mode:

1. Enter `dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t lssi`
2. Wait for the command to process and display the end results.

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.

- 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 tucs0n -script ~/bin/mkpprcpairs`
2. Wait for the script to process and provide a report regarding the success or failure of the process.

Here is an example 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 four DS CLI commands that you have used. It also allows you to repeat any of the last four 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.

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

   Note: If you should make a mistake and type the wrong user name or password, do not try to correct this within the current session. Exit the DS CLI session 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.

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>! a list of commands is displayed 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 the third last command.

Activating licenses

After you have installed your storage unit and DS Storage Manager, your first step is to activate your licenses.

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

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

The initial enablement of any optional 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:

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

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

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

1. Start the DS Storage Manager application.
2. In the navigation panel select **Real-time Manager** → **Manage Hardware**, then select **Storage Units**.

3. In the storage unit main page select the storage unit, click **Properties** in the **Select Action** drop-down list, and then click **Go**. The properties page displays for the storage unit.
4. Gather the following information about your storage unit. You must enter this information at the IBM Web site in the next step of this task. You can use the following table to document this information.
   - From the **MTMS** fill in the information 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's serial number. For example, if the MTMS is IBM.1750.511.75FA120, then the machine type is 1750, the model number is 511, and the machine serial number is 75FA120.
   - From the **Machine signature** field, note the machine signature.

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

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

   Note: In most situations, the DSFA application can locate your 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 will need the OCC (which you can find on the License Function Authorization document).

Importing activation codes:

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

Note: Before you begin this task, you must resolve any current DS6000 problems. Refer to the IBM TotalStorage DS6000 Troubleshooting information. If you need additional assistance to resolve these problems, contact IBM Support.

1. In the navigation panel, under Manage Hardware, select Storage units. In Storage units — Main Page, select a Storage unit. Then select the **Activation codes** tab. The Activation codes page is displayed.
2. Click **Import key file**. The Import page is displayed.
3. In the **Select file to import** field, specify the target file. Use the **Browse** button to navigate to the appropriate directory.
4. After you have specified the code file, click **OK** to complete the process.

Applying activation codes:

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

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

Note: Before you begin this task, you must resolve any current DS6000 problems. Refer to the *IBM TotalStorage DS6000 Installation, Troubleshooting, and Recovery Guide*. If you need additional assistance to resolve these problems, contact IBM Support.

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

1. In the navigation panel, select, in order, Real-time Manager, Manage Hardware, and storage units. In Storage units — Main Page, select a storage unit. Then select **Configure** in the **Select Action** drop-down list, and then click **Go**. Select 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 contained within the activation codes.
2. Click **Apply**, then **OK** to complete the process.

Defining Multiple Management Consoles

Use this process to create a storage complex domain by establishing a connection with a secondary Management Console for redundancy.

This process must be done from the primary Management Console. You must have the Management Console IP address and the appropriate user ID and password.

When you are adding a peer Management Console, the peer Management Console cannot have storage units that are associated with it. If you have two Management Consoles that each have storage units that are associated with them, you must remove the storage units from the Management Console that you want to use as the peer Management Console. After you remove the storage units, add the peer Management Console to the primary Management Console. You can then add the previously removed storage units through the primary Management Console.

This task enables you to identify and establish a connection with a secondary Management Console for the Storage complex associated with the primary Management Console.

1. Under **Manage hardware**, select **Storage complexes**. In Storage Complexes — Main Page, select **Define peer** in the **Select Action** drop-down box. Then click **Go**. The Define peer management console page is displayed.

2. Specify the server host name for the secondary Management Console.
3. Specify the server IP address.
4. Click **Ok**. The storage complex domain is established.

Specifying storage unit day and time

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

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

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

Reviewing storage unit network settings

Use this process 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.

1. In the navigation, select **Real-time Manager**, **Manage Hardware**, and then **Storage units**. Select the appropriate storage unit.
2. In the **Select Action** drop-down list, select **Configure** and then **Go**. The Storage unit properties page is displayed.
3. In the navigation on the left, click **Network settings**. The Network settings tab 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.

Defining customer contacts

Use this process to add or modify shipping or contact information for a customer account.

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 **Customer contact** and then **Go**. The Customer account information tab is displayed. From this page, you can create or modify the customer account information.
4. In the navigation on the left, click Shipping information to add or modify shipping information for the customer.
5. In the navigation on the left, click Contact information to add or modify contact information for the customer.
6. Click the **OK** button to complete the customer contact information.

Configuring I/O ports

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

1. In the navigation panel, under Manage Hardware, select Storage units. In Storage units — Main Page, select a Storage unit. Select **Configure I/O Ports...** in the **Select Action** drop-down list. Then click **Go**. The Configure I/O Ports page is displayed.
2. Use the check boxes to select one or more host attachments of the same type.
3. In the **Select Action** drop-down, select the I/O port type that you want to change to. You can change any I/O port to FcAl, FcSf, or FICON. Then click **Go**. The table will update with the attachment type that you selected.

Creating host systems

Use this process to create host systems and define their parameters.

You must have at least one array and one rank defined before creating hosts.

1. In the navigation, under Manage Hardware, select Host Systems. In Host Systems — Main Page, select a storage complex (and possibly a storage unit), and select **Create...** in the **Select Action** drop-down list. 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. Otherwise, go to step 5.
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 Defined 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, specify the storage units for the host attachment field by selecting a storage unit from the Available storage units list and click **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 (You can optionally choose **Select volume group later** if you do not want to select the volume group now), and choose a login option. You can loop through this page for each host attachment identifier by selecting the **Apply assignment** button to commit the current transaction and then starting from the top by selecting 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 login to **the following specific storage unit I/O ports**, then 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 **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 then specify the correct values. Otherwise, click **Finish** to complete the host system creation process.

Creating open systems volumes

Use this process to create open systems volumes and to specify their attributes and properties.

1. In the navigation, under Configure Storage, select Open Systems. Under Open Systems, select Volumes — Open Systems. In Volumes — Open Systems, open the **Select Action** drop-down list and select **Create...**. Then click **Go**. The Create Volume — Select extent pool page is displayed.
2. You must select an extent pool for the target volumes. You can use the **Create new extent pool** button to create a new extent pool. After you create the extent pool and the table resets, the new extent pool is available for selection.
3. After you select the extent pool, click **Next**. The Create Volume — Define volume characteristics page is displayed.
4. Define the characteristics for the target volumes. You can select any number of volume groups from the **Select volume groups** list to associate with the target volumes.
   a. Select the volume type, volume groups, and optionally select the **Enable write cache with mirroring** selection.
   b. Optionally, use the **Create new group** button (selected by default) to create a new volume group.
5. After you define the volume characteristics, click **Next**. The Create Volume — Define volume properties page is displayed.
6. Define the volume properties. Use the **Calculate max quantity** button to populate the **Quantity** field with the calculated value. To calculate the maximum size, enter a value in the **quantity** field and click the **Calculate max size** button to see the maximum size. To calculate the maximum quantity, enter a value in the **Size** field and click the **Calculate max quantity** button to see the maximum quantity.
7. If you select the **Calculate max size** button, the **Size** field is populated with the calculated value. You can overwrite the value and enter a lesser size, and you can enter only an integer. 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.
8. 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. The **Available storage in extent pool** field displays the amount of usable storage in this extent pool in GB.
9. After you define the volume properties, click **Next**. The Create Volume — Create volume nicknames page is displayed.
10. Optionally, you can 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.

11. In Verification, review the attributes and values to verify that they are correct.
12. 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 volume creation process.

Creating iSeries volumes

Use this process to quickly configure a storage complex with iSeries volumes.

1. Under **Configure Storage**, select **Express Configuration Wizard**. The Express configuration wizard page is displayed.
2. Select the storage unit for the volumes that you are configuring under **Select storage unit**.
3. Select **iSeries (FB)** from the **Select volume type** list, 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 following to configure an amount of available storage:
   - Select **Amount of unused storage to configure** 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.
7. Click **Calculate**. The value for the third item is automatically calculated.
8. To enable host creation for the iSeries volumes, select **Create host**.
9. Click **Next**. The Set volume naming page is displayed.
10. To specify a sequence of volume names for the quantity created, select **Generate a sequence of nicknames based on the following**.
11. Select **Verify nicknames are unique** to determine whether volume names that are generated on this page are unique.
12. Enter the prefix (alphabetic) that you want for the volumes in the **Prefix** box.
13. Enter the suffix (numeric) that you want for the volumes in the **Numeric suffix** box, and then click **Next**.
14. 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 then click **Next**. The General host information page appears.
    - If you did not select to create a host, the Verification page appears and you can go to the last step.
    - If you selected to create a host, the General host information page appears and you can continue to the next step.
15. 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.
16. Optionally, enter a description (up to 256 characters) for the host, and then click **Next**. The Host ports page appears.

17. 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.
18. 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 then click **Next**. The Define WWPNs page appears.
19. 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 groups page is displayed.
20. To map the volume group to the host attachment, select **Assign host attachment to volume group**, and then click **Next**.
21. 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

Use this process to quickly and easily configure a storage complex with zSeries volumes.

1. Under **Configure storage**, select **Express configuration wizard**. The Express configuration wizard page is displayed.
2. Select the storage unit for the volumes that you are configuring under **Select storage unit**.
3. Select **zSeries (CKD)** from the **Select volume type** list, 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 to configure an amount of available storage:
   - Select **Amount of unused storage to configure**. 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:
    - 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, 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 a real-time storage configuration

Use this process to create a real-time storage configuration that includes zSeries volumes or open systems volumes.

You must have your storage unit installed. See “Creating a storage unit (Simulated only)” on page 65 for more information.

This process covers one method for creating a configuration on the storage unit. Your network must be configured so that it can support all the components and functions that you will use with your storage unit.

1. Assign the storage unit to a storage complex. In the navigation, under **Manage Hardware**, select **Storage Complexes**. Select a storage complex from the table. See “Assigning a storage unit to a storage complex” on page 66 for more information.

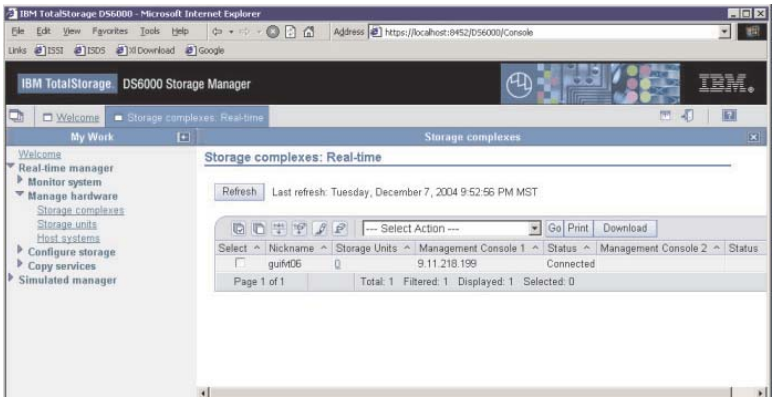

f2d00100

Note: At this point, you can now install and use the DSCLI to complete the remaining post installation tasks, or you can continue using the DS Storage manager. See “Using the DS CLI application” on page 68 for more information.

2. Activate licenses. In the navigation, under **Real-time manager**, select **Manage Hardware**, then select **Storage Units**. See “Activating licenses” on page 71 for more information.
3. (Optional) Define a peer storage complex by clicking on **Storage complexes** and selecting a storage complex from the table. See “Defining Multiple Management Consoles” on page 73 for more information.
   a. Click **Define Peer** in the **Select Action** drop-down list to assign the storage unit to the selected storage complex, and then click **Go**.

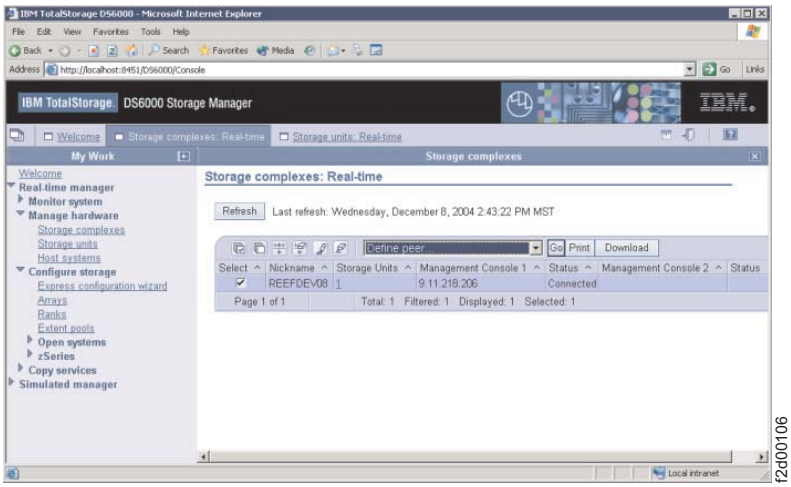

f2d00106

b. Define either the hostname or IP address (not both) for the peer management console. Click **OK** when you are finished.

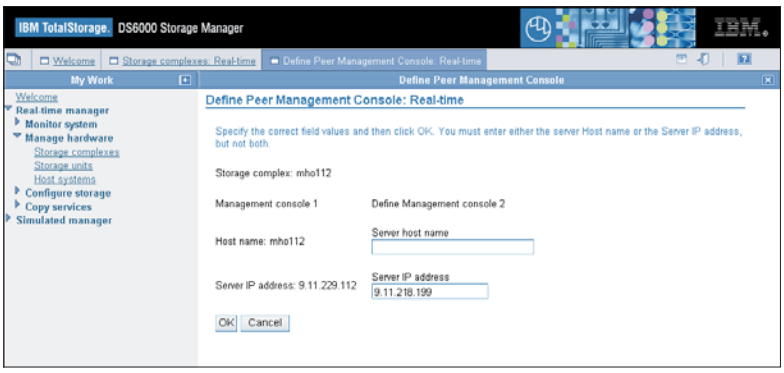

f2d00107

c. Click on the **Storage Complexes** main page and you should see a second management console nickname defined with status on the right most columns. See “Assigning a storage unit to a storage complex” on page 66 for more information.

4. Configure the storage unit. In the navigation, select **Manage Hardware**, then **Storage Units**. Select a storage unit in the table. In the **Select Action** drop-down box, select **Configure** and then **Go**. See “Specifying storage unit day and time” on page 74 and “Reviewing storage unit network settings” on page 74 for more information.
5. Define customer contact information. In the navigation, select **Manage Hardware**, then **Storage Units**. In the **Select Action** drop-down box, select **Customer contact** and then **Go**. See “Defining customer contacts” on page 74 for more information.
6. Configure notifications. In the navigation, select **Manage Hardware**, then **Storage Units**. In the **Select Action** drop-down box, select **Configure notifications** and then **Go**. See “Setting up call home” on page 123 for more information.
7. Configure the I/O ports. In the navigation panel, under **Manage Hardware**, select **Storage Units**. In Storage units — Main Page, select a Storage unit. Select **Configure I/O Ports...** in the **Select Action** drop-down list. Then click **Go**. See “Configuring I/O ports” on page 74 for more information.
8. Create an array. In the navigation, select **Real-time Manager** → **Configure storage** → **Arrays**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating arrays” on page 91 for more information.

9. Create a rank. In the navigation, select **Real-time Manager** → **Configure storage** → **Ranks**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating ranks” on page 92 for more information.
10. Create the host system. In the navigation, select **Manage Hardware**, then **Host Systems**. In the **Select Action** drop-down box, select **Create** and then **Go**. See “Creating host systems” on page 75 for more information.
11. Create an extent pool. In the navigation, select **Real-time Manager** → **Configure storage** → **Extent pools**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating extent pools” on page 93 for more information.
12. Perform one of the following tasks:
    - Create zSeries LCUs and zSeries volumes.
      a. Create zSeries LCUs. In the navigation, select **Real-time Manager** → **Configure storage** → **zSeries** → **LCUs**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating LCUs” on page 94 for more information.
      b. Create zSeries volumes. In the navigation, select **Real-time Manager** → **Configure storage** → **zSeries** → **Volumes**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating zSeries volumes” on page 78 for more information.
    - Create open systems volumes and open systems volume groups.
      a. Create open systems volumes. In the navigation, select **Real-time Manager** → **Configure storage** → **Open systems** → **Volumes - Open Systems**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating open systems volumes” on page 76 for more information.
      b. Create open systems volume groups. In the navigation, select **Real-time Manager** → **Configure storage** → **Open systems** → **Volume groups**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating open systems volume groups” on page 96 for more information.

Creating a storage unit (Simulated only)

Use this process to create a storage unit and to specify its attributes and properties.

1. In the navigation, select, in order, Simulated Manager, Manage Hardware, and Storage Units. In Storage Unit — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create Storage Unit — General storage unit information page is displayed.
2. 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 need 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.
3. Click **Next** to continue.
4. The Create Storage Unit —Specify DDM packs page is displayed. You must specify the **Quantity of DDM packs** and the **DDM type**. Click **Add**, then click **Next** to continue. The Define licensed function page is displayed.
5. In the Define licensed function page, you must specify a value in the **Operating Environment License (TB)** field.

a. 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.

6. Specify values in the remaining four fields as appropriate. Click **Next** to continue.
7. The Verification page is displayed. Use this page to review the established 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 storage unit creation process.

Assigning a storage unit to a storage complex

Use this process 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.

1. In the navigation, select, in order, Real-time Manager, Manage Hardware, and Storage Complexes. In Storage Complexes— Main Page, select the appropriate storage complex from the table.

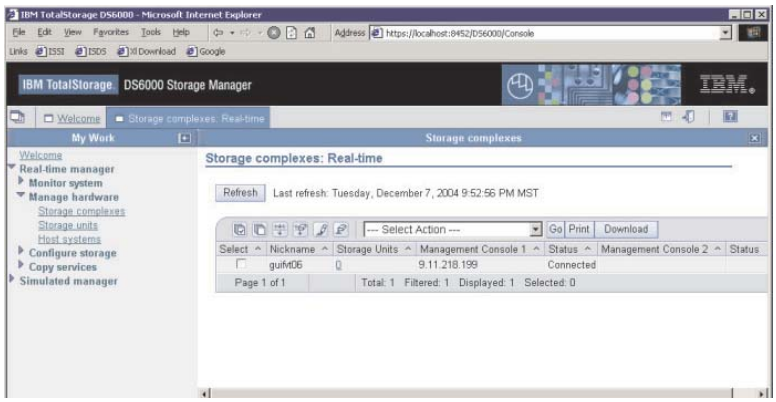

f2d00100

2. In the **Select Action** drop-down list, select **Assign Storage unit** and then **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 serial number. The serial number must be at least 7 digits, with the first 2 digits being the point of manufacturer (for example, if the serial number is 13AAXRA, the point of manufacturer is 13).

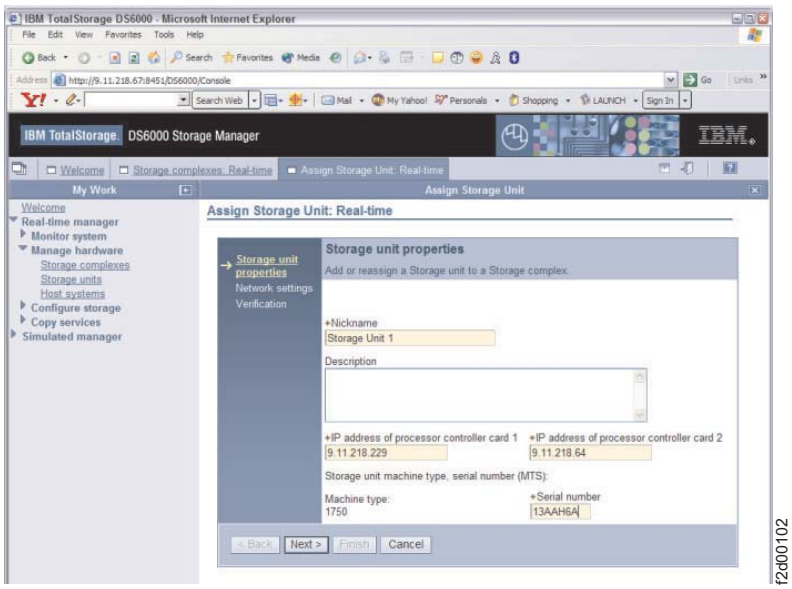

f2d00102

3. Click **Next**. The Network settings page is displayed.

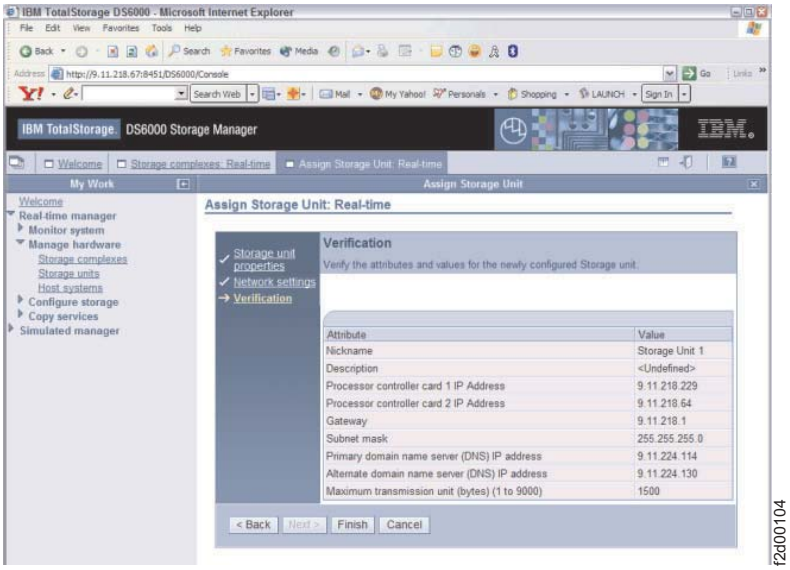

f2d00104

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. Enter a different Max transmission units value, if necessary.

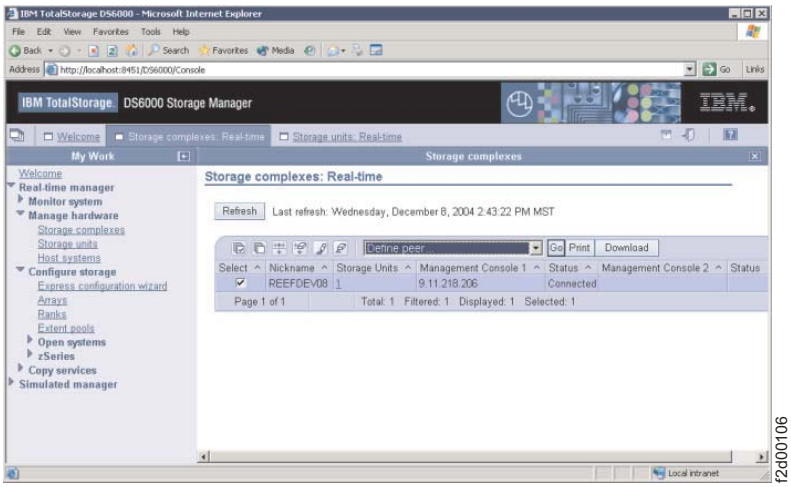

f2d00106

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

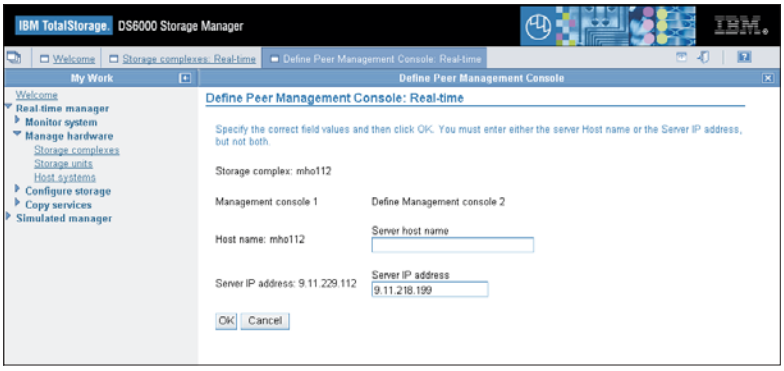

f2d00107

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, there are three command modes that are available to you.

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

- Single-shot
- Interactive
- Script

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

Sun Solaris
/opt/ibm/dscli

Windows
C:\Program Files\IBM\dscli

HP Tru64
/opt/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 is expired when 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`.
- Supply the log in information in a profile configuration file (for additional information, see the topic ″Default configuration setup with a profile file″). 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.

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. Use the following example to use the single-shot mode:

1. Enter `dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t lssi`
2. Wait for the command to process and display the end results.

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.

- 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 tucs0n -script ~/bin/mkpprcpairs`
2. Wait for the script to process and provide a report regarding the success or failure of the process.

Here is an example 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 four DS CLI commands that you have used. It also allows you to repeat any of the last four 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.

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

   Note: If you should make a mistake and type the wrong user name or password, do not try to correct this within the current session. Exit the DS CLI session 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.

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>! a list of commands is displayed 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 the third last command.

Activating licenses

After you have installed your storage unit and DS Storage Manager, your first step is to activate your licenses.

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

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

The initial enablement of any optional 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:

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

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

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

1. Start the DS Storage Manager application.
2. In the navigation panel select **Real-time Manager** → **Manage Hardware**, then select **Storage Units**.

3. In the storage unit main page select the storage unit, click **Properties** in the **Select Action** drop-down list, and then click **Go**. The properties page displays for the storage unit.
4. Gather the following information about your storage unit. You must enter this information at the IBM Web site in the next step of this task. You can use the following table to document this information.
   - From the **MTMS** fill in the information 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's serial number. For example, if the MTMS is IBM.1750.511.75FA120, then the machine type is 1750, the model number is 511, and the machine serial number is 75FA120.
   - From the **Machine signature** field, note the machine signature.

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

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

   Note: In most situations, the DSFA application can locate your 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 will need the OCC (which you can find on the License Function Authorization document).

Importing activation codes:

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

Note: Before you begin this task, you must resolve any current DS6000 problems. Refer to the IBM TotalStorage DS6000 Troubleshooting information. If you need additional assistance to resolve these problems, contact IBM Support.

1. In the navigation panel, under Manage Hardware, select Storage units. In Storage units — Main Page, select a Storage unit. Then select the **Activation codes** tab. The Activation codes page is displayed.
2. Click **Import key file**. The Import page is displayed.
3. In the **Select file to import** field, specify the target file. Use the **Browse** button to navigate to the appropriate directory.
4. After you have specified the code file, click **OK** to complete the process.

Applying activation codes:

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

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

Note: Before you begin this task, you must resolve any current DS6000 problems. Refer to the *IBM TotalStorage DS6000 Installation, Troubleshooting, and Recovery Guide*. If you need additional assistance to resolve these problems, contact IBM Support.

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

1. In the navigation panel, select, in order, Real-time Manager, Manage Hardware, and storage units. In Storage units — Main Page, select a storage unit. Then select **Configure** in the **Select Action** drop-down list, and then click **Go**. Select 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 contained within the activation codes.
2. Click **Apply**, then **OK** to complete the process.

Defining Multiple Management Consoles

Use this process to create a storage complex domain by establishing a connection with a secondary Management Console for redundancy.

This process must be done from the primary Management Console. You must have the Management Console IP address and the appropriate user ID and password.

When you are adding a peer Management Console, the peer Management Console cannot have storage units that are associated with it. If you have two Management Consoles that each have storage units that are associated with them, you must remove the storage units from the Management Console that you want to use as the peer Management Console. After you remove the storage units, add the peer Management Console to the primary Management Console. You can then add the previously removed storage units through the primary Management Console.

This task enables you to identify and establish a connection with a secondary Management Console for the Storage complex associated with the primary Management Console.

1. Under **Manage hardware**, select **Storage complexes**. In Storage Complexes — Main Page, select **Define peer** in the **Select Action** drop-down box. Then click **Go**. The Define peer management console page is displayed.

2. Specify the server host name for the secondary Management Console.
3. Specify the server IP address.
4. Click **Ok**. The storage complex domain is established.

Specifying storage unit day and time

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

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

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

Reviewing storage unit network settings

Use this process 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.

1. In the navigation, select **Real-time Manager**, **Manage Hardware**, and then **Storage units**. Select the appropriate storage unit.
2. In the **Select Action** drop-down list, select **Configure** and then **Go**. The Storage unit properties page is displayed.
3. In the navigation on the left, click **Network settings**. The Network settings tab 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.

Defining customer contacts

Use this process to add or modify shipping or contact information for a customer account.

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 **Customer contact** and then **Go**. The Customer account information tab is displayed. From this page, you can create or modify the customer account information.
4. In the navigation on the left, click Shipping information to add or modify shipping information for the customer.
5. In the navigation on the left, click Contact information to add or modify contact information for the customer.
6. Click the **OK** button to complete the customer contact information.

Configuring I/O ports

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

1. In the navigation panel, under Manage Hardware, select Storage units. In Storage units — Main Page, select a Storage unit. Select **Configure I/O Ports...** in the **Select Action** drop-down list. Then click **Go**. The Configure I/O Ports page is displayed.
2. Use the check boxes to select one or more host attachments of the same type.
3. In the **Select Action** drop-down, select the I/O port type that you want to change to. You can change any I/O port to FcAl, FcSf, or FICON. Then click **Go**. The table will update with the attachment type that you selected.

Creating arrays

Use this process to create arrays, either automatic or custom, and to specify their RAID and rank attributes.

1. In the navigation, under Configure Storage, select Arrays. In Arrays — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create Array — Definition method page is displayed.
2. The Definition method step provides you with the option of specifying the array sites yourself or having the application specify them. The choice you make here determines the next step in the process of creating arrays.
   a. Choose **Create arrays automatically** to specify the quantity and RAID type. Click **Next** and go to step 3, or,
   b. Choose **Create custom arrays** to select the RAID type and the array site numbers. Click **Next** and go to step 4. (see note at c.)
   c. There is also a checkbox for **Create an 8 disk array** that changes the following steps. If a & c are selected, 2 arrays, each with 2 array-sites, are created. If b & c are selected, and you only want to create one array, then you must select only one array-site on the Array Configuration (custom) page. If b & c are selected, and you want 2 arrays, select 2 array-sites and click Next. You will then need to select a secondary array-site for each of the two primary array-sites (as described in step 5).
3. In Create Array — Array configuration (Auto), specify the quantity and RAID type for the arrays. Then click **Next**. The Add array to rank page is displayed. Go to step 6.
4. In Create Array — Array configuration (custom), select from the list of supported RAID types and select at least one array site. To create a second array-site, select the **Create an 8 disk array** checkbox. Then click **Next**. If you selected **Create an 8 disk array** in Create Array — Array configuration (custom), the Create Array — Second array-site selection page is displayed. Go to step 5. If you did not select **Create an 8 disk array**, the Create Array — Add array to rank page is displayed. Go to step 6.
5. In Create Array — Second array-site selection, for each array site for this configuration, the compatible array-sites appear for you to select a second array-site. Select the appropriate second array-sites. Then click **Next**. The Create Array — Add array to rank page is displayed. Go to step 6.
6. In Create Array — Add array to rank, a check in the box for **Add these arrays to ranks** specifies that the new arrays are to be put into ranks. Uncheck the box if you do not want to put the new arrays into ranks. If this box is checked, you must select a value in the **Storage type** field. Then click **Next**. The Create Array — Verification page is displayed.
7. In Create Array — Verification, review the attributes and values to verify that they are correct.
8. 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.

Creating ranks

Use this process to create ranks and to specify their arrays and extent pools.

1. In the navigation, under Configure Storage, select Ranks. In Ranks — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create Rank — Select array for rank page is displayed.
2. In the **Select** column, select an array. Then click **Next**. The Create Rank — Define rank properties page is displayed.
3. The **Rank number** is provided by default. Specify the **Storage type** and then click **Next**. The Create Rank — Select extent pool page is displayed.
4. Optionally, you can select one extent pool. Either make a selection in the **Select** column, or click on the **Create extent pool** button.

   Note: If you select the **Create extent pool** button, follow the process for creating a new extent pool. Once you have completed that process by clicking on the **Finish** button, the new extent pool is available for selection in the **Select** column.
5. After you have selected an extent pool, click on **Next**. The Create Rank — Verification page is displayed.
6. In Create Rank — Verification, review the attributes and values to verify that they are correct.
7. 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.

Creating host systems

Use this process to create host systems and define their parameters.

You must have at least one array and one rank defined before creating hosts.

1. In the navigation, under Manage Hardware, select Host Systems. In Host Systems — Main Page, select a storage complex (and possibly a storage unit), and select **Create...** in the **Select Action** drop-down list. 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. Otherwise, go to step 5.
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 Defined 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, specify the storage units for the host attachment field by selecting a storage unit from the Available storage units list and click **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 (You can optionally choose **Select volume group later** if you do not want to select the volume group now), and choose a login option. You can loop through this page for each host attachment identifier by selecting the **Apply assignment** button to commit the current transaction and then starting from the top by selecting 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 login to **the following specific storage unit I/O ports**, then 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 **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 then specify the correct values. Otherwise, click **Finish** to complete the host system creation process.

Creating extent pools

Use this process to create extent pools, automatic or custom, and to specify extent pool parameters.

1. In the navigation, under Configure Storage, select Extent Pools. In Extent Pools — Main Page, open the **Select Action** drop-down list and select **Create...**. Then click **Go**. The Create Extent Pool — Definition method page is displayed.
2. The Definition method step provides you with the option of having the necessary arrays and ranks automatically created or of selecting the ranks yourself for the extent pool.
   a. Choose **Create extent pool automatically based on storage requirements . . .** to have the arrays and ranks automatically created and put into the extent pool. Then click **Next** and go to step 3, or,
   b. Choose **Create custom extent pool . . .** to select the ranks for the extent pool. Then click **Next** and go to step 4.
3. You chose to create the extent pool automatically. The Create Extent Pool — Define extent pool requirements page is displayed. Specify nickname, storage type, RAID type, and the required amount of storage. Check the box for **Use any existing unassigned arrays and ranks** to use arrays or ranks that were already created. Click **Next** and go to step 5 to define the reserve storage.
4. You chose to create a custom extent pool. The Create Extent Pool — Define extent pool properties page is displayed. Specify nickname, storage type, RAID type, and server. Click **Next** and go to step 6.
5. The Create Extent Pool — Reserve storage page is displayed. Specify the percent of reserved storage in the extent pool. Click **Next** and the Create Extent Pool — Verification page is displayed. Go to step 8.
6. The Create Extent Pool — Create ranks page is displayed. You must select at least one rank.
   a. If you select the **Create new rank** button, follow the process for creating a new rank. Once you have completed that process by clicking the **Finish** button, the new rank is available for selection in the **Select** column.

7. After you have selected the ranks, click **Next**. The Create Extent Pool — Verification page is displayed.
8. In Verification, 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 then specify the correct values. Otherwise, click **Finish** to complete the extent pool creation process.

Creating LCUs

Use this process to create logical control units (LCUs) and to specify their attributes and properties.

1. In the navigation, select Configure Storage, zSeries, and LCUs. In LCUs — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create LCU — Select from available LCUs page is displayed.
2. In the Select from available LCUs page, select one or more LCUs from the list of those available. **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 associated with Server 1 can be associated only with LCUs that have odd numbers.
3. Click **Next** to continue. The Define LCU properties page is displayed.
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.
7. 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 LCU creation process.

Creating zSeries volumes

Use this process to quickly and easily configure a storage complex with zSeries volumes.

1. Under **Configure storage**, select **Express configuration wizard**. The Express configuration wizard page is displayed.
2. Select the storage unit for the volumes that you are configuring under **Select storage unit**.
3. Select **zSeries (CKD)** from the **Select volume type** list, 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 to configure an amount of available storage:
   - Select **Amount of unused storage to configure**. 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:
    - 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, 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 open systems volumes

Use this process to create open systems volumes and to specify their attributes and properties.

1. In the navigation, under Configure Storage, select Open Systems. Under Open Systems, select Volumes — Open Systems. In Volumes — Open Systems, open the **Select Action** drop-down list and select **Create...**. Then click **Go**. The Create Volume — Select extent pool page is displayed.
2. You must select an extent pool for the target volumes. You can use the **Create new extent pool** button to create a new extent pool. After you create the extent pool and the table resets, the new extent pool is available for selection.
3. After you select the extent pool, click **Next**. The Create Volume — Define volume characteristics page is displayed.
4. Define the characteristics for the target volumes. You can select any number of volume groups from the **Select volume groups** list to associate with the target volumes.
    a. Select the volume type, volume groups, and optionally select the **Enable write cache with mirroring** selection.
    b. Optionally, use the **Create new group** button (selected by default) to create a new volume group.
5. After you define the volume characteristics, click **Next**. The Create Volume — Define volume properties page is displayed.
6. Define the volume properties. Use the **Calculate max quantity** button to populate the **Quantity** field with the calculated value. To calculate the maximum size, enter a value in the **quantity** field and click the **Calculate max**

size button to see the maximum size. To calculate the maximum quantity, enter a value in the **Size** field and click the **Calculate max quantity** button to see the maximum quantity.

7. If you select the **Calculate max size** button, the **Size** field is populated with the calculated value. You can overwrite the value and enter a lesser size, and you can enter only an integer. 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.
8. 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. The **Available storage in extent pool** field displays the amount of usable storage in this extent pool in GB.
9. After you define the volume properties, click **Next**. The Create Volume — Create volume nicknames page is displayed.
10. Optionally, you can 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.
11. In Verification, review the attributes and values to verify that they are correct.
12. 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 volume creation process.

Creating open systems volume groups

Use this process to create open systems volume groups and to specify their properties and parameters.

1. In the navigation, under Configure Storage, select Open Systems. Under Open Systems, select Volume Groups. In Volume Groups — Main Page, open the **Select Action** drop-down list and select **Create...**. Then click **Go**. The Create Volume Group — Define volume group properties page is displayed.
2. Define the properties. Both the **Nickname** and **Accessed by host types** fields are required. Select the appropriate host type. All valid host types will be automatically selected.
3. After you define the properties, click **Next**. The Create Volume Group — Select host attachments page is displayed.
4. Optionally, you can select the host attachment for the target volumes. You can use the **Create new host attachment** button to create a new host attachment. After you have created the new host attachment, the table resets and the new host attachment is available for selection.
5. After selecting the host attachment, click **Next**. The Create Volume Group — Select volumes for groups page is displayed.

6. 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. You can use the **Create new volumes** button to create a new volume. After you create the volume and the table resets, the new volume is 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.

   Note: If an unexpected error occurs, you must click **Cancel** to exit the wizard.
7. After selecting the volumes for the volume group, click **Next**. The Create Volume Group — Verification page is displayed.
8. Use the Verification page to review the attributes and verify that they are correct.
9. 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 volume group creation process.

Creating simulated configurations

The topics in this section provide information for creating simulated storage configurations.

Creating an express simulated storage configuration

Use this process to create an express simulated storage configuration.

You must have your storage unit installed. See “Creating a storage unit (Simulated only)” on page 65 for more information.

This scenario covers one method for creating a configuration on the storage unit. Your network must be configured so that it can support all the components and functions that you will use with your storage unit.

1. Create a configuration file. In the navigation, under **Simulated Manager**, select **Manage Configuration Files**. Click **Create new** in the **Select Action** drop-down list, and then click **Go**. See “Creating configuration files (Simulated only)” on page 98 for more information.
2. Create the storage unit. In the navigation, under **Simulated Manager**, select **Manage Hardware**, then **Storage Units**. Click **Create** in the **Select Action** drop-down list, and then click **Go**. See “Creating a storage unit (Simulated only)” on page 65 for more information.
3. Assign the storage unit to a storage complex. In the navigation, under **Manage Hardware**, select **Storage Complexes**. Select a storage complex from the table. See “Assigning a storage unit to a storage complex” on page 66 for more information.

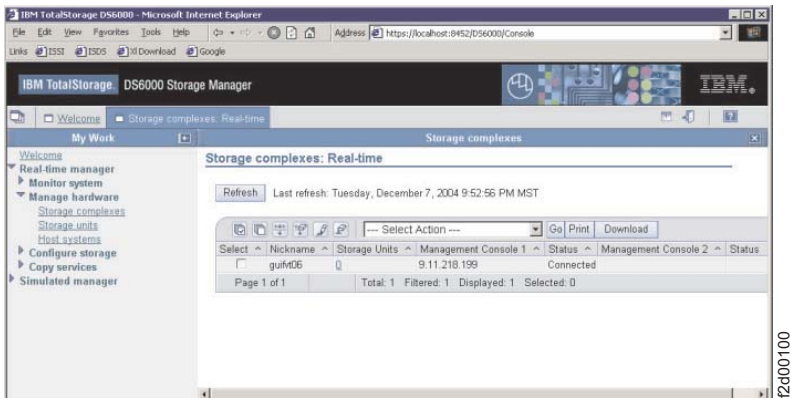

Note: At this point, you can now install and use the DSCLI to complete the remaining post installation tasks, or you can continue using the DS Storage manager. See “Using the DS CLI application” on page 68 for more information.

4. Define customer contact information. In the navigation, select **Manage Hardware**, then **Storage Units**. In the **Select Action** drop-down box, select **Customer contact** and then **Go**. See “Defining customer contacts” on page 74 for more information.
5. Configure notifications. In the navigation, select **Manage Hardware**, then **Storage Units**. In the **Select Action** drop-down box, select **Configure notifications** and then **Go**. See “Setting up call home” on page 123 for more information.
6. Create the host system. In the navigation, select **Manage Hardware**, then **Host Systems**. In the **Select Action** drop-down box, select **Create** and then **Go**. See “Creating host systems” on page 75 for more information.
7. Configure I/O ports for the host system. In the navigation, select **Simulated Manager** → **Manage hardware** → **Storage units**. Select the storage unit to configure. In the Select Action drop-down box, select **Configure I/O Ports...** and then **Go**. See “Configuring I/O ports” on page 74 for more information.
8. Configure storage. In the navigation, select **Simulated Manager**, then **Configure Storage**, then **Express Configuration Wizard**. See “Creating open systems volumes” on page 76, “Creating iSeries volumes” on page 77, and “Creating zSeries volumes” on page 78 for more information.

   Note: The express configuration tool defines and configures for only one host, not multiple hosts.
9. Apply the configuration. In the navigation, select **Manage Hardware**, then **Storage Units**. In the **Select Action** drop-down box, select **Apply configuration** and then **Go**. See “Applying a configuration (Simulated only)” on page 109 for more information.

Creating configuration files (Simulated only)

Use this process to create a simulated configuration file.

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 **Manage Configuration Files**. Select **Create new...** from the **Section Action** drop-down list to create a new enterprise file for offline configuration.

2. If you have another file open when you select the **Create new...** action, a message prompts you to save your current work before creating a new file.
3. You must either import a storage unit instance, or create a new storage unit instance from the Create storage unit wizard pages.
4. This enterprise file 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.
5. Use the **Save** action to complete the process.

Creating a storage unit (Simulated only)

Use this process to create a storage unit and to specify its attributes and properties.

1. In the navigation, select, in order, Simulated Manager, Manage Hardware, and Storage Units. In Storage Unit — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create Storage Unit — General storage unit information page is displayed.
2. 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 need 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.
3. Click **Next** to continue.
4. The Create Storage Unit —Specify DDM packs page is displayed. You must specify the **Quantity of DDM packs** and the **DDM type**. Click **Add**, then click **Next** to continue. The Define licensed function page is displayed.
5. In the Define licensed function page, you must specify a value in the **Operating Environment License (TB)** field.
   a. 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.
6. Specify values in the remaining four fields as appropriate. Click **Next** to continue.
7. The Verification page is displayed. Use this page to review the established 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 storage unit creation process.

Assigning a storage unit to a storage complex

Use this process 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.

1. In the navigation, select, in order, Real-time Manager, Manage Hardware, and Storage Complexes. In Storage Complexes— Main Page, select the appropriate storage complex from the table.

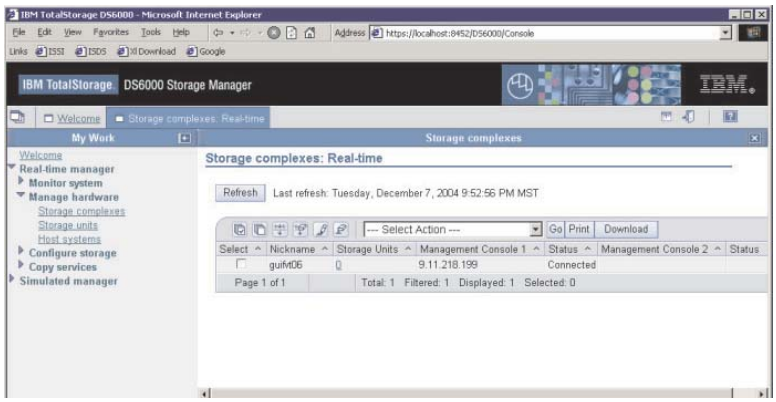

f2d00100

2. In the **Select Action** drop-down list, select **Assign Storage unit** and then **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 serial number. The serial number must be at least 7 digits, with the first 2 digits being the point of manufacturer (for example, if the serial number is 13AAXRA, the point of manufacturer is 13).

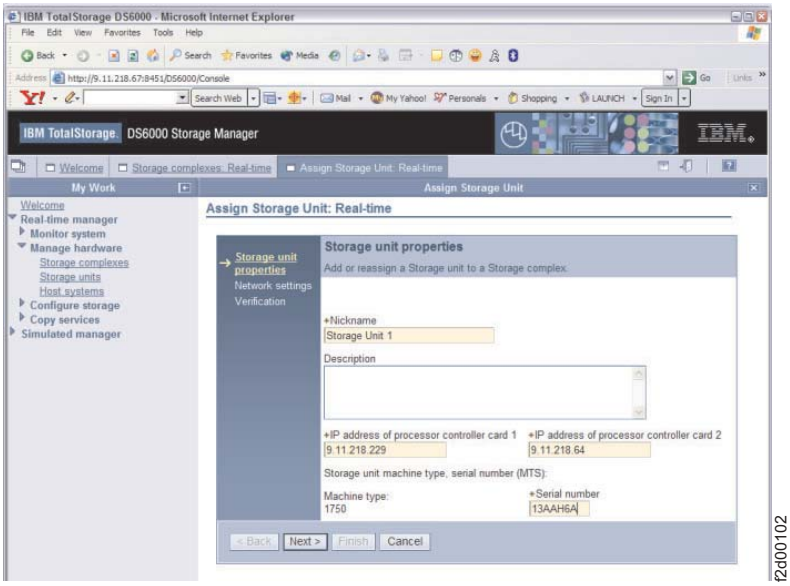

f2d00102

3. Click **Next**. The Network settings page is displayed.

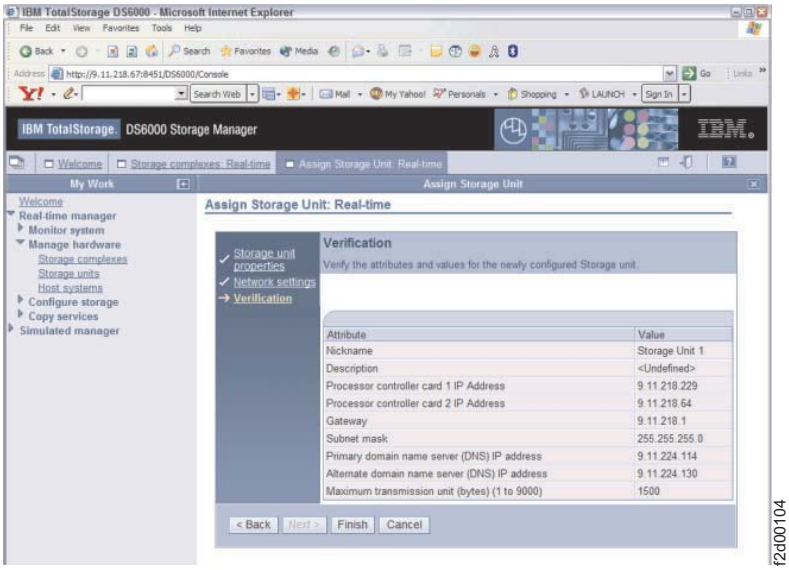

4. Specify the appropriate network settings and then click **Next**. The Verification page is displayed.
   a. Enter a gateway
   b. Enter a subnet mask
   c. Enter the primary DNS address
   d. Enter the secondary DNS address
   e. Enter a different Max transmission units value, if necessary.

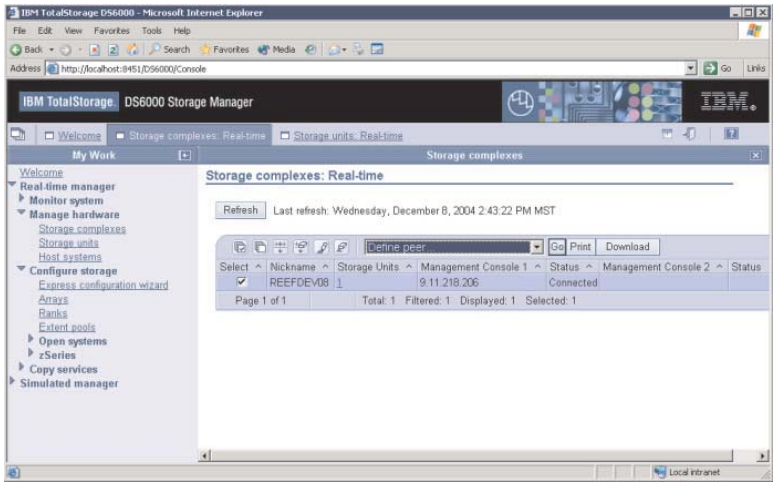

5. Verify the attributes and values for the newly configured Storage unit. Click **Finish** if the settings are correct.

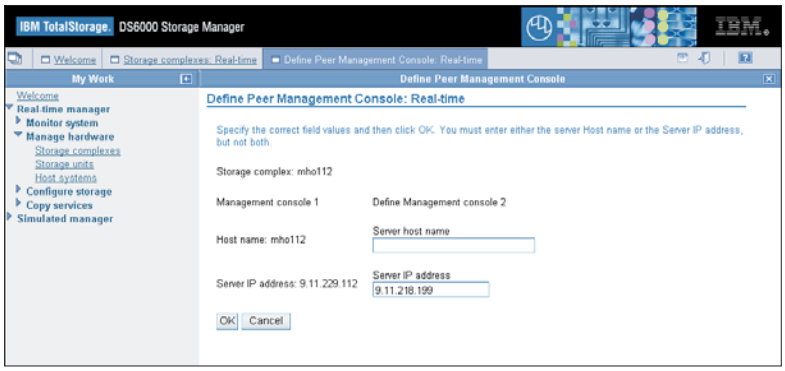

f2d00107

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, there are three command modes that are available to you.

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

- Single-shot
- Interactive
- Script

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

Sun Solaris
/opt/ibm/dscli

Windows
C:\Program Files\IBM\dscli

HP Tru64
/opt/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 is expired when 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`.
- Supply the log in information in a profile configuration file (for additional information, see the topic ″Default configuration setup with a profile file″). 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.

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. Use the following example to use the single-shot mode:

1. Enter `dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t lssi`
2. Wait for the command to process and display the end results.

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.

- 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 tucs0n -script ~/bin/mkpprcpairs`
2. Wait for the script to process and provide a report regarding the success or failure of the process.

Here is an example 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 four DS CLI commands that you have used. It also allows you to repeat any of the last four 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.

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

   Note: If you should make a mistake and type the wrong user name or password, do not try to correct this within the current session. Exit the DS CLI session 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.

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>! a list of commands is displayed 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 the third last command.

Defining customer contacts

Use this process to add or modify shipping or contact information for a customer account.

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 **Customer contact** and then **Go**. The Customer account information tab is displayed. From this page, you can create or modify the customer account information.
4. In the navigation on the left, click Shipping information to add or modify shipping information for the customer.
5. In the navigation on the left, click Contact information to add or modify contact information for the customer.
6. Click the **OK** button to complete the customer contact information.

Configuring I/O ports

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

1. In the navigation panel, under Manage Hardware, select Storage units. In Storage units — Main Page, select a Storage unit. Select **Configure I/O Ports...** in the **Select Action** drop-down list. Then click **Go**. The Configure I/O Ports page is displayed.
2. Use the check boxes to select one or more host attachments of the same type.
3. In the **Select Action** drop-down, select the I/O port type that you want to change to. You can change any I/O port to FcAl, FcSf, or FICON. Then click **Go**. The table will update with the attachment type that you selected.

Creating host systems

Use this process to create host systems and define their parameters.

You must have at least one array and one rank defined before creating hosts.

1. In the navigation, under Manage Hardware, select Host Systems. In Host Systems — Main Page, select a storage complex (and possibly a storage unit), and select **Create...** in the **Select Action** drop-down list. 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. Otherwise, go to step 5.
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 Defined 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, specify the storage units for the host attachment field by selecting a storage unit from the Available storage units list and click **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 (You can optionally choose **Select volume group later** if you do not want to select the volume group now), and choose a login option. You can loop through this page for each host attachment identifier by selecting the **Apply assignment** button to commit the current transaction and then starting from the top by selecting 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 login to **the following specific storage unit I/O ports**, then 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 **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 then specify the correct values. Otherwise, click **Finish** to complete the host system creation process.

Creating open systems volumes

Use this process to create open systems volumes and to specify their attributes and properties.

1. In the navigation, under Configure Storage, select Open Systems. Under Open Systems, select Volumes — Open Systems. In Volumes — Open Systems, open the **Select Action** drop-down list and select **Create...**. Then click **Go**. The Create Volume — Select extent pool page is displayed.
2. You must select an extent pool for the target volumes. You can use the **Create new extent pool** button to create a new extent pool. After you create the extent pool and the table resets, the new extent pool is available for selection.
3. After you select the extent pool, click **Next**. The Create Volume — Define volume characteristics page is displayed.
4. Define the characteristics for the target volumes. You can select any number of volume groups from the **Select volume groups** list to associate with the target volumes.
   a. Select the volume type, volume groups, and optionally select the **Enable write cache with mirroring** selection.
   b. Optionally, use the **Create new group** button (selected by default) to create a new volume group.
5. After you define the volume characteristics, click **Next**. The Create Volume — Define volume properties page is displayed.
6. Define the volume properties. Use the **Calculate max quantity** button to populate the **Quantity** field with the calculated value. To calculate the maximum size, enter a value in the **quantity** field and click the **Calculate max**

size button to see the maximum size. To calculate the maximum quantity, enter a value in the **Size** field and click the **Calculate max quantity** button to see the maximum quantity.

7. If you select the **Calculate max size** button, the **Size** field is populated with the calculated value. You can overwrite the value and enter a lesser size, and you can enter only an integer. 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.
8. 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. The **Available storage in extent pool** field displays the amount of usable storage in this extent pool in GB.
9. After you define the volume properties, click **Next**. The Create Volume — Create volume nicknames page is displayed.
10. Optionally, you can 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.
11. In Verification, review the attributes and values to verify that they are correct.
12. 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 volume creation process.

Creating iSeries volumes

Use this process to quickly configure a storage complex with iSeries volumes.

1. Under **Configure Storage**, select **Express Configuration Wizard**. The Express configuration wizard page is displayed.
2. Select the storage unit for the volumes that you are configuring under **Select storage unit**.
3. Select **iSeries (FB)** from the **Select volume type** list, 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 following to configure an amount of available storage:
   - Select **Amount of unused storage to configure** 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.
7. Click **Calculate**. The value for the third item is automatically calculated.

8. To enable host creation for the iSeries volumes, select **Create host**.
9. Click **Next**. The Set volume naming page is displayed.
10. To specify a sequence of volume names for the quantity created, select **Generate a sequence of nicknames based on the following**.
11. Select **Verify nicknames are unique** to determine whether volume names that are generated on this page are unique.
12. Enter the prefix (alphabetic) that you want for the volumes in the **Prefix** box.
13. Enter the suffix (numeric) that you want for the volumes in the **Numeric suffix** box, and then click **Next**.
14. 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 then click **Next**. The General host information page appears.
    - If you did not select to create a host, the Verification page appears and you can go to the last step.
    - If you selected to create a host, the General host information page appears and you can continue to the next step.
15. 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.
16. Optionally, enter a description (up to 256 characters) for the host, and then click **Next**. The Host ports page appears.
17. 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.
18. 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 then click **Next**. The Define WWPNs page appears.
19. 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 groups page is displayed.
20. To map the volume group to the host attachment, select **Assign host attachment to volume group**, and then click **Next**.
21. 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

Use this process to quickly and easily configure a storage complex with zSeries volumes.

1. Under **Configure storage**, select **Express configuration wizard**. The Express configuration wizard page is displayed.
2. Select the storage unit for the volumes that you are configuring under **Select storage unit**.
3. Select **zSeries (CKD)** from the **Select volume type** list, 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 to configure an amount of available storage:
    - Select **Amount of unused storage to configure**. 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:
    - 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, 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**.

Applying a configuration (Simulated only)

Use this process to select, authenticate, and apply a storage unit configuration.

1. In the navigation, under Manage hardware, select Storage units. In Storage units — Main Page, open the **Select Action** drop-down list and select **Apply configuration...**. Then click **Go**. The Apply Configuration — Select application method page is displayed.
2. 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 network connected 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 Authentication page is displayed. Click **Next** to continue and go to step 4.

3. Use the Select storage complex page to connect directly to a storage unit. Click **Next** to continue. The Authentication page is displayed.
4. In the Authentication 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.
5. In the Select storage unit page, you connect directly to a storage unit. Specify the values as appropriate and click **Next**.
6. Use the Verification page to review the attributes and verify that they are correct.
7. 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.

Creating a simulated storage configuration

Use this process to create a simulated storage configuration that includes zSeries volumes or open systems volumes.

You must have your storage unit installed. See “Creating a storage unit (Simulated only)” on page 65 for more information.

This process covers one method for creating a configuration on the storage unit. Your network must be configured so that it can support all the components and functions that you will use with your storage unit.

1. Create a configuration file. In the navigation, under **Simulated Manager**, select **Manage Configuration Files**. Click **Create new** in the **Select Action** drop-down list, and then click **Go**. See “Creating configuration files (Simulated only)” on page 98 for more information.
2. Create a storage complex. In the navigation, select **Simulated Manager** → **Manage hardware** → **Storage complexes**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating a storage complex (Simulated only)” on page 113 for more information.
3. Create the storage unit. In the navigation, under **Simulated Manager**, select **Manage Hardware**, then **Storage Units**. Click **Create** in the **Select Action** drop-down list, and then click **Go**. See “Creating a storage unit (Simulated only)” on page 65 for more information.
4. Assign the storage unit to a storage complex. In the navigation, under **Manage Hardware**, select **Storage Complexes**. Select a storage complex from the table. See “Assigning a storage unit to a storage complex” on page 66 for more information.

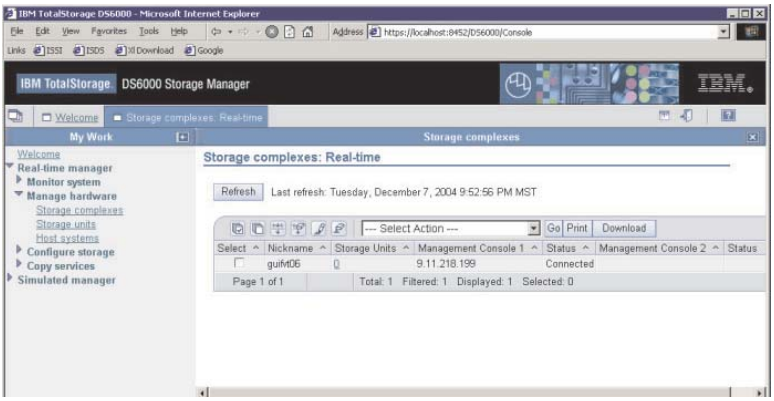

f2d00100

Note: At this point, you can now install and use the DSCLI to complete the remaining post installation tasks, or you can continue using the DS Storage manager. See “Using the DS CLI application” on page 68 for more information.

5. Activate licenses. In the navigation, under **Simulated manager**, select **Manage Hardware**, then select **Storage Units**. See “Activating licenses” on page 71 for more information.
6. (Optional) Define a peer storage complex by clicking on **Storage complexes** and selecting a storage complex from the table. See “Defining Multiple Management Consoles” on page 73 for more information.
   a. Click **Define Peer** in the **Select Action** drop-down list to assign the storage unit to the selected storage complex, and then click **Go**.

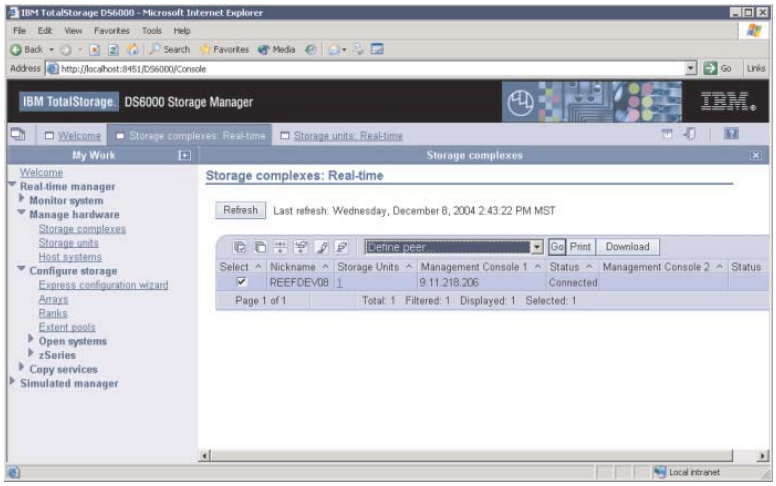

f2d00106

   b. Define either the hostname or IP address (not both) for the peer management console. Click **OK** when you are finished.

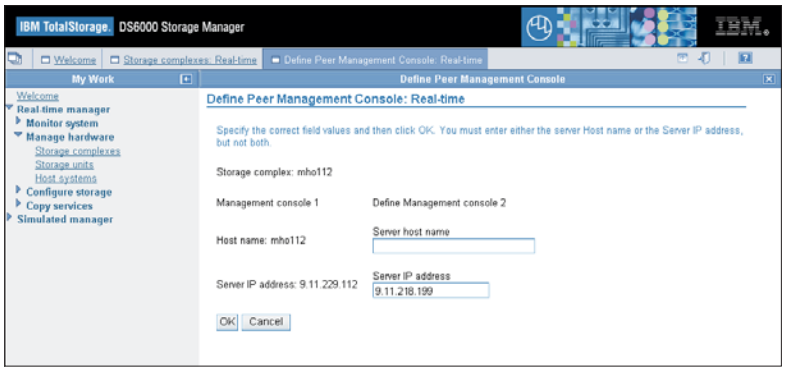

f2d00107

   c. Click on the **Storage Complexes** main page and you should see a second management console nickname defined with status on the right most columns. See “Assigning a storage unit to a storage complex” on page 66 for more information.
7. Configure the storage unit. In the navigation, select **Manage Hardware**, then **Storage Units**. Select a storage unit in the table. In the **Select Action** drop-down box, select **Configure** and then **Go**. See “Specifying storage unit day and time” on page 74 and “Reviewing storage unit network settings” on page 74 for more information.

8. Define customer contact information. In the navigation, select **Manage Hardware**, then **Storage Units**. In the **Select Action** drop-down box, select **Customer contact** and then **Go**. See “Defining customer contacts” on page 74 for more information.
9. Configure notifications. In the navigation, select **Manage Hardware**, then **Storage Units**. In the **Select Action** drop-down box, select **Configure notifications** and then **Go**. See “Setting up call home” on page 123 for more information.
10. Configure the I/O ports. In the navigation panel, under **Manage Hardware**, select **Storage Units**. In Storage units — Main Page, select a Storage unit. Select **Configure I/O Ports...** in the **Select Action** drop-down list. Then click **Go**. See “Configuring I/O ports” on page 74 for more information.
11. Create the host system. In the navigation, select **Manage Hardware**, then **Host Systems**. In the **Select Action** drop-down box, select **Create** and then **Go**. See “Creating host systems” on page 75 for more information.
12. Create an array. In the navigation, select **Simulated Manager** → **Configure storage** → **Arrays**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating arrays” on page 91 for more information.
13. Create a rank. In the navigation, select **Simulated Manager** → **Configure storage** → **Ranks**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating ranks” on page 92 for more information.
14. Create an extent pool. In the navigation, select **Simulated Manager** → **Configure storage** → **Extent pools**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating extent pools” on page 93 for more information.
15. Perform one of the following tasks:
    - Create zSeries LCUs and zSeries volumes.
      a. Create zSeries LCUs. In the navigation, select **Simulated Manager** → **Configure storage** → **zSeries** → **LCUs**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating LCUs” on page 94 for more information.
      b. Create zSeries volumes. In the navigation, select **Simulated Manager** → **Configure storage** → **zSeries** → **Volumes**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating zSeries volumes” on page 78 for more information.
    - Create open systems volumes and open systems volume groups.
      a. Create open systems volumes. In the navigation, select **Simulated Manager** → **Configure storage** → **Open systems** → **Volumes - Open Systems**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating open systems volumes” on page 76 for more information.
      b. Create open systems volume groups. In the navigation, select **Simulated Manager** → **Configure storage** → **Open systems** → **Volume groups**. In the Select Action drop-down box, select **Create...** and then **Go**. See “Creating open systems volume groups” on page 96 for more information.
16. Apply the configuration. In the navigation, select **Manage Hardware**, then **Storage Units**. In the **Select Action** drop-down box, select **Apply configuration** and then **Go**. See “Applying a configuration (Simulated only)” on page 109 for more information.

Creating a storage complex (Simulated only)

Use this process to create a storage complex and to specify its nickname and storage unit assignments.

1. In the navigation, select, in order, Simulated Manager, Manage Hardware, and Storage Complexes. In Storage Complexes— Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create Storage Complex — Define properties page is displayed.
2. 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 downloading or uploading 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.
3. After you have defined the properties, click **Next** to continue. The Verification page is displayed.
4. Use the Verification page to review the established attributes and verify that they are correct.
5. 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.

Creating a storage unit (Simulated only)

Use this process to create a storage unit and to specify its attributes and properties.

1. In the navigation, select, in order, Simulated Manager, Manage Hardware, and Storage Units. In Storage Unit — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create Storage Unit — General storage unit information page is displayed.
2. 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 need 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.
3. Click **Next** to continue.
4. The Create Storage Unit —Specify DDM packs page is displayed. You must specify the **Quantity of DDM packs** and the **DDM type**. Click **Add**, then click **Next** to continue. The Define licensed function page is displayed.
5. In the Define licensed function page, you must specify a value in the **Operating Environment License (TB)** field.
   a. 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.
6. Specify values in the remaining four fields as appropriate. Click **Next** to continue.
7. The Verification page is displayed. Use this page to review the established 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 storage unit creation process.

Creating configuration files (Simulated only)

Use this process to create a simulated configuration file.

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 **Manage Configuration Files**. Select **Create new...** from the **Section Action** drop-down list to create a new enterprise file for offline configuration.
2. If you have another file open when you select the **Create new...** action, a message prompts you to save your current work before creating a new file.
3. You must either import a storage unit instance, or create a new storage unit instance from the Create storage unit wizard pages.
4. This enterprise file 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.
5. Use the **Save** action to complete the process.

Creating a storage unit (Simulated only)

Use this process to create a storage unit and to specify its attributes and properties.

1. In the navigation, select, in order, Simulated Manager, Manage Hardware, and Storage Units. In Storage Unit — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create Storage Unit — General storage unit information page is displayed.
2. 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 need 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.
3. Click **Next** to continue.
4. The Create Storage Unit —Specify DDM packs page is displayed. You must specify the **Quantity of DDM packs** and the **DDM type**. Click **Add**, then click **Next** to continue. The Define licensed function page is displayed.
5. In the Define licensed function page, you must specify a value in the **Operating Environment License (TB)** field.
   a. 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.
6. Specify values in the remaining four fields as appropriate. Click **Next** to continue.
7. The Verification page is displayed. Use this page to review the established 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 storage unit creation process.

Creating a storage unit (Simulated only)

Use this process to create a storage unit and to specify its attributes and properties.

1. In the navigation, select, in order, Simulated Manager, Manage Hardware, and Storage Units. In Storage Unit — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create Storage Unit — General storage unit information page is displayed.
2. 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 need 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.
3. Click **Next** to continue.
4. The Create Storage Unit —Specify DDM packs page is displayed. You must specify the **Quantity of DDM packs** and the **DDM type**. Click **Add**, then click **Next** to continue. The Define licensed function page is displayed.
5. In the Define licensed function page, you must specify a value in the **Operating Environment License (TB)** field.
   a. 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.
6. Specify values in the remaining four fields as appropriate. Click **Next** to continue.
7. The Verification page is displayed. Use this page to review the established 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 storage unit creation process.

Assigning a storage unit to a storage complex

Use this process 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.

1. In the navigation, select, in order, Real-time Manager, Manage Hardware, and Storage Complexes. In Storage Complexes— Main Page, select the appropriate storage complex from the table.

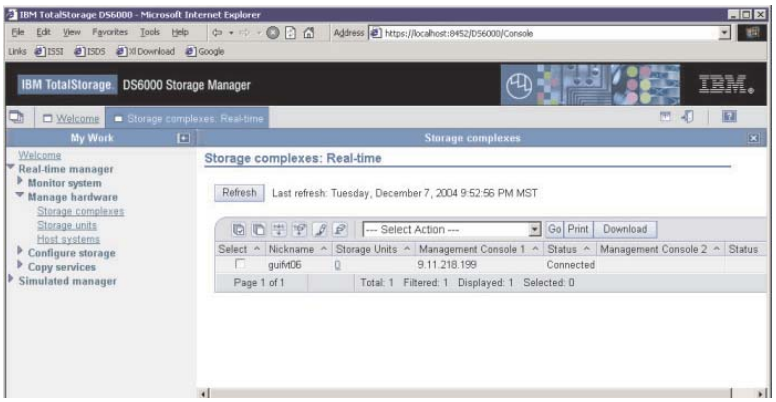

f2d00100

2. In the **Select Action** drop-down list, select **Assign Storage unit** and then **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 serial number. The serial number must be at least 7 digits, with the first 2 digits being the point of manufacturer (for example, if the serial number is 13AAXRA, the point of manufacturer is 13).

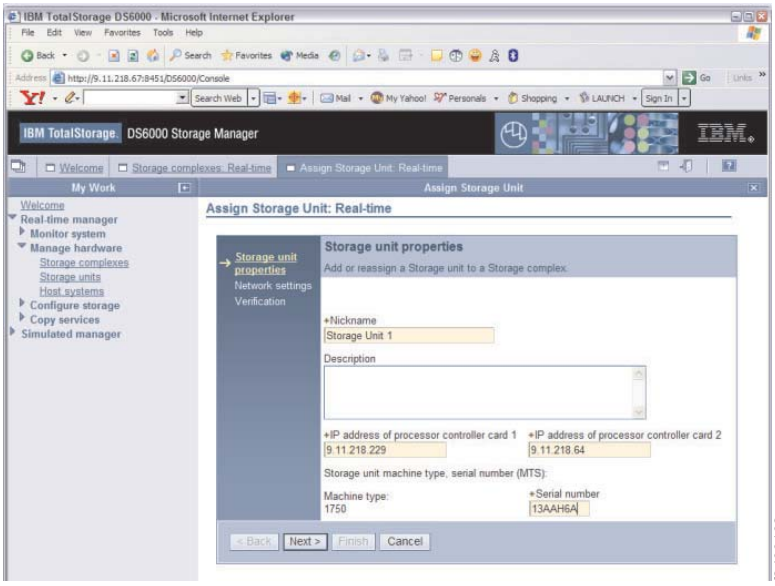

3. Click **Next**. The Network settings page is displayed.

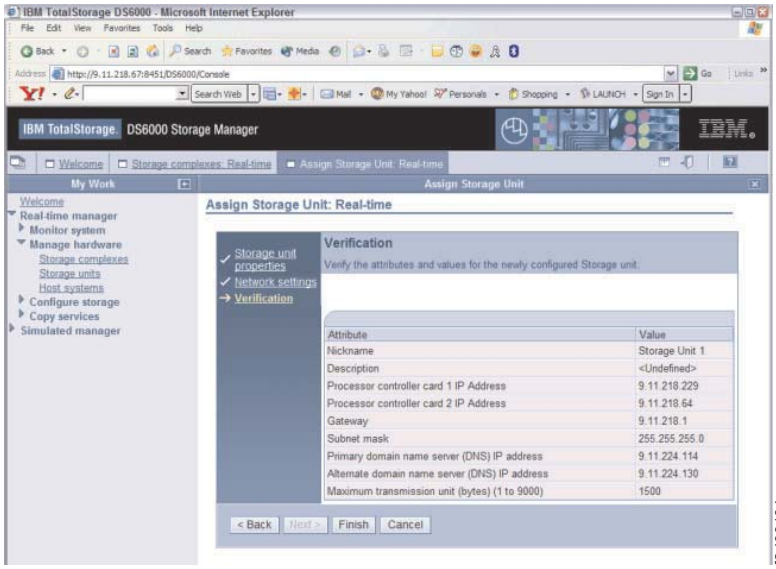

4. Specify the appropriate network settings and then click **Next**. The Verification page is displayed.

   a. Enter a gateway

   b. Enter a subnet mask

   c. Enter the primary DNS address

   d. Enter the secondary DNS address

   e. Enter a different Max transmission units value, if necessary.

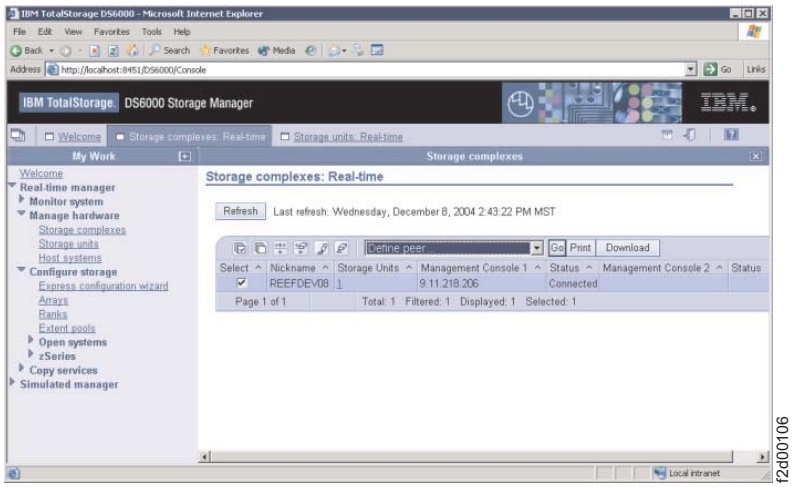

f2d00106

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

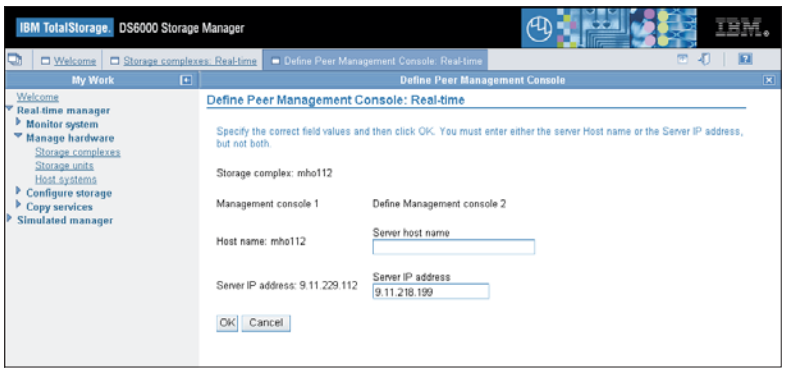

f2d00107

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, there are three command modes that are available to you.

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

- Single-shot
- Interactive
- Script

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

Sun Solaris
/opt/ibm/dscli

Windows
C:\Program Files\IBM\dscli

HP Tru64
/opt/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 is expired when 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`.
- Supply the log in information in a profile configuration file (for additional information, see the topic ″Default configuration setup with a profile file″). 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.

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. Use the following example to use the single-shot mode:

1. Enter `dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t lssi`
2. Wait for the command to process and display the end results.

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.

- 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 tucs0n -script ~/bin/mkpprcpairs`
2. Wait for the script to process and provide a report regarding the success or failure of the process.

Here is an example 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 four DS CLI commands that you have used. It also allows you to repeat any of the last four 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.

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

   Note: If you should make a mistake and type the wrong user name or password, do not try to correct this within the current session. Exit the DS CLI session 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.

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>! a list of commands is displayed 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 the third last command.

Activating licenses

After you have installed your storage unit and DS Storage Manager, your first step is to activate your licenses.

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

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

The initial enablement of any optional 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:

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

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

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

1. Start the DS Storage Manager application.
2. In the navigation panel select **Real-time Manager** → **Manage Hardware**, then select **Storage Units**.

3. In the storage unit main page select the storage unit, click **Properties** in the **Select Action** drop-down list, and then click **Go**. The properties page displays for the storage unit.
4. Gather the following information about your storage unit. You must enter this information at the IBM Web site in the next step of this task. You can use the following table to document this information.
   - From the **MTMS** fill in the information 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's serial number. For example, if the MTMS is IBM.1750.511.75FA120, then the machine type is 1750, the model number is 511, and the machine serial number is 75FA120.
   - From the **Machine signature** field, note the machine signature.

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

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

   Note: In most situations, the DSFA application can locate your 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 will need the OCC (which you can find on the License Function Authorization document).

Importing activation codes:

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

Note: Before you begin this task, you must resolve any current DS6000 problems. Refer to the IBM TotalStorage DS6000 Troubleshooting information. If you need additional assistance to resolve these problems, contact IBM Support.

1. In the navigation panel, under Manage Hardware, select Storage units. In Storage units — Main Page, select a Storage unit. Then select the **Activation codes** tab. The Activation codes page is displayed.
2. Click **Import key file**. The Import page is displayed.
3. In the **Select file to import** field, specify the target file. Use the **Browse** button to navigate to the appropriate directory.
4. After you have specified the code file, click **OK** to complete the process.

Applying activation codes:

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

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

Note: Before you begin this task, you must resolve any current DS6000 problems. Refer to the *IBM TotalStorage DS6000 Installation, Troubleshooting, and Recovery Guide*. If you need additional assistance to resolve these problems, contact IBM Support.

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

1. In the navigation panel, select, in order, Real-time Manager, Manage Hardware, and storage units. In Storage units — Main Page, select a storage unit. Then select **Configure** in the **Select Action** drop-down list, and then click **Go**. Select 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 contained within the activation codes.
2. Click **Apply**, then **OK** to complete the process.

Defining Multiple Management Consoles

Use this process to create a storage complex domain by establishing a connection with a secondary Management Console for redundancy.

This process must be done from the primary Management Console. You must have the Management Console IP address and the appropriate user ID and password.

When you are adding a peer Management Console, the peer Management Console cannot have storage units that are associated with it. If you have two Management Consoles that each have storage units that are associated with them, you must remove the storage units from the Management Console that you want to use as the peer Management Console. After you remove the storage units, add the peer Management Console to the primary Management Console. You can then add the previously removed storage units through the primary Management Console.

This task enables you to identify and establish a connection with a secondary Management Console for the Storage complex associated with the primary Management Console.

1. Under **Manage hardware**, select **Storage complexes**. In Storage Complexes — Main Page, select **Define peer** in the **Select Action** drop-down box. Then click **Go**. The Define peer management console page is displayed.

2. Specify the server host name for the secondary Management Console.
3. Specify the server IP address.
4. Click **Ok**. The storage complex domain is established.

Specifying storage unit day and time

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

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

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

Reviewing storage unit network settings

Use this process 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.

1. In the navigation, select **Real-time Manager**, **Manage Hardware**, and then **Storage units**. Select the appropriate storage unit.
2. In the **Select Action** drop-down list, select **Configure** and then **Go**. The Storage unit properties page is displayed.
3. In the navigation on the left, click **Network settings**. The Network settings tab 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.

Setting up call home

Use this process to configure notifications.

You must define the customer contact information before configuring notifications.

This task enables you to define Call Home (SMTP), SNMP, and SIMs (service information messages) notifications for a storage unit. See Notification methods for detailed information about these functions.

1. Under **Manage hardware**, select **Storage units**. In Storage units - Main Page, select **Configure notifications** in the **Select Action** drop-down field. Then click **Go**. The Configure notifications - Define Call Home page is displayed.
2. Ensure that **Enable Call Home** is selected to activate Call Home. (This is checked by default.)
3. Complete the SMTP information.
   a. Enter the SMTP server host name.
   b. Enter the SMTP IP address.
   c. Enter the SMTP server ports.
   d. Click **Apply**.

4. Click **Test Call Home connection** to send a connection test and generate a problem log entry. A confirmation message is displayed.
5. Click the **SNMP** tab. The Define SNMP connection page is displayed.
6. Select Enable SNMP notification to define the SNMP connection properties for the selected storage units.
7. Specify either an IP address, a Host name, or both under SNMP trap destination.
8. Specify an SNMP community name of up to 32 characters. This field is used to authenticate requests. 'Public' is selected by default.
9. Optionally specify an SNMP system contact name of up to 32 characters. Enter a destination port. Click **Apply**.
10. Click the **zSeries** tab. The Define Service Information Messages for zSeries page is displayed.
11. Optionally select a SIM severity level in the Severity reporting level for DASD Service Information Messages field.
12. Optionally select a Media Service Information severity level in the Severity reporting level for Media Service Information Messages field.
13. Optionally select a Service Information severity level in the Severity reporting level for a Service Information Messages field.
14. Click Apply.

Defining customer contacts

Use this process to add or modify shipping or contact information for a customer account.

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 **Customer contact** and then **Go**. The Customer account information tab is displayed. From this page, you can create or modify the customer account information.
4. In the navigation on the left, click Shipping information to add or modify shipping information for the customer.
5. In the navigation on the left, click Contact information to add or modify contact information for the customer.
6. Click the **OK** button to complete the customer contact information.

Configuring I/O ports

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

1. In the navigation panel, under Manage Hardware, select Storage units. In Storage units — Main Page, select a Storage unit. Select **Configure I/O Ports...** in the **Select Action** drop-down list. Then click **Go**. The Configure I/O Ports page is displayed.
2. Use the check boxes to select one or more host attachments of the same type.
3. In the **Select Action** drop-down, select the I/O port type that you want to change to. You can change any I/O port to FcAl, FcSf, or FICON. Then click **Go**. The table will update with the attachment type that you selected.

Creating arrays

Use this process to create arrays, either automatic or custom, and to specify their RAID and rank attributes.

1. In the navigation, under Configure Storage, select Arrays. In Arrays — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create Array — Definition method page is displayed.
2. The Definition method step provides you with the option of specifying the array sites yourself or having the application specify them. The choice you make here determines the next step in the process of creating arrays.
   a. Choose **Create arrays automatically** to specify the quantity and RAID type. Click **Next** and go to step 3, or,
   b. Choose **Create custom arrays** to select the RAID type and the array site numbers. Click **Next** and go to step 4. (see note at c.)
   c. There is also a checkbox for **Create an 8 disk array** that changes the following steps. If a & c are selected, 2 arrays, each with 2 array-sites, are created. If b & c are selected, and you only want to create one array, then you must select only one array-site on the Array Configuration (custom) page. If b & c are selected, and you want 2 arrays, select 2 array-sites and click Next. You will then need to select a secondary array-site for each of the two primary array-sites (as described in step 5).
3. In Create Array — Array configuration (Auto), specify the quantity and RAID type for the arrays. Then click **Next**. The Add array to rank page is displayed. Go to step 6.
4. In Create Array — Array configuration (custom), select from the list of supported RAID types and select at least one array site. To create a second array-site, select the **Create an 8 disk array** checkbox. Then click **Next**. If you selected **Create an 8 disk array** in Create Array — Array configuration (custom), the Create Array — Second array-site selection page is displayed. Go to step 5. If you did not select **Create an 8 disk array**, the Create Array — Add array to rank page is displayed. Go to step 6.
5. In Create Array — Second array-site selection, for each array site for this configuration, the compatible array-sites appear for you to select a second array-site. Select the appropriate second array-sites. Then click **Next**. The Create Array — Add array to rank page is displayed. Go to step 6.
6. In Create Array — Add array to rank, a check in the box for **Add these arrays to ranks** specifies that the new arrays are to be put into ranks. Uncheck the box if you do not want to put the new arrays into ranks. If this box is checked, you must select a value in the **Storage type** field. Then click **Next**. The Create Array — Verification page is displayed.
7. In Create Array — Verification, review the attributes and values to verify that they are correct.
8. 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.

Creating ranks

Use this process to create ranks and to specify their arrays and extent pools.

1. In the navigation, under Configure Storage, select Ranks. In Ranks — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create Rank — Select array for rank page is displayed.
2. In the **Select** column, select an array. Then click **Next**. The Create Rank — Define rank properties page is displayed.
3. The **Rank number** is provided by default. Specify the **Storage type** and then click **Next**. The Create Rank — Select extent pool page is displayed.
4. Optionally, you can select one extent pool. Either make a selection in the **Select** column, or click on the **Create extent pool** button.

Note: If you select the **Create extent pool** button, follow the process for creating a new extent pool. Once you have completed that process by clicking on the **Finish** button, the new extent pool is available for selection in the **Select** column.

5. After you have selected an extent pool, click on **Next**. The Create Rank — Verification page is displayed.
6. In Create Rank — Verification, review the attributes and values to verify that they are correct.
7. 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.

Creating host systems

Use this process to create host systems and define their parameters.

You must have at least one array and one rank defined before creating hosts.

1. In the navigation, under Manage Hardware, select Host Systems. In Host Systems — Main Page, select a storage complex (and possibly a storage unit), and select **Create...** in the **Select Action** drop-down list. 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. Otherwise, go to step 5.
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 Defined 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, specify the storage units for the host attachment field by selecting a storage unit from the Available storage units list and click **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 (You can optionally choose **Select volume group later** if you do not want to select the volume group now), and choose a login option. You can loop through this page for each host attachment identifier by selecting the **Apply assignment** button to commit the current transaction and then starting from the top by selecting 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 login to **the following specific storage unit I/O ports**, then 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 **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 then specify the correct values. Otherwise, click **Finish** to complete the host system creation process.

Creating extent pools

Use this process to create extent pools, automatic or custom, and to specify extent pool parameters.

1. In the navigation, under Configure Storage, select Extent Pools. In Extent Pools — Main Page, open the **Select Action** drop-down list and select **Create...**. Then click **Go**. The Create Extent Pool — Definition method page is displayed.
2. The Definition method step provides you with the option of having the necessary arrays and ranks automatically created or of selecting the ranks yourself for the extent pool.
   a. Choose **Create extent pool automatically based on storage requirements . . .** to have the arrays and ranks automatically created and put into the extent pool. Then click **Next** and go to step 3, or,
   b. Choose **Create custom extent pool . . .** to select the ranks for the extent pool. Then click **Next** and go to step 4.
3. You chose to create the extent pool automatically. The Create Extent Pool — Define extent pool requirements page is displayed. Specify nickname, storage type, RAID type, and the required amount of storage. Check the box for **Use any existing unassigned arrays and ranks** to use arrays or ranks that were already created. Click **Next** and go to step 5 to define the reserve storage.
4. You chose to create a custom extent pool. The Create Extent Pool — Define extent pool properties page is displayed. Specify nickname, storage type, RAID type, and server. Click **Next** and go to step 6.
5. The Create Extent Pool — Reserve storage page is displayed. Specify the percent of reserved storage in the extent pool. Click **Next** and the Create Extent Pool — Verification page is displayed. Go to step 8.
6. The Create Extent Pool — Create ranks page is displayed. You must select at least one rank.
   a. If you select the **Create new rank** button, follow the process for creating a new rank. Once you have completed that process by clicking the **Finish** button, the new rank is available for selection in the **Select** column.
7. After you have selected the ranks, click **Next**. The Create Extent Pool — Verification page is displayed.
8. In Verification, 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 then specify the correct values. Otherwise, click **Finish** to complete the extent pool creation process.

Creating LCUs

Use this process to create logical control units (LCUs) and to specify their attributes and properties.

1. In the navigation, select Configure Storage, zSeries, and LCUs. In LCUs — Main Page, select **Create...** in the **Select Action** drop-down list. Then click **Go**. The Create LCU — Select from available LCUs page is displayed.

2. In the Select from available LCUs page, select one or more LCUs from the list of those available. **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 associated with Server 1 can be associated only with LCUs that have odd numbers.
3. Click **Next** to continue. The Define LCU properties page is displayed.
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.
7. 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 LCU creation process.

Creating zSeries volumes

Use this process to quickly and easily configure a storage complex with zSeries volumes.

1. Under **Configure storage**, select **Express configuration wizard**. The Express configuration wizard page is displayed.
2. Select the storage unit for the volumes that you are configuring under **Select storage unit**.
3. Select **zSeries (CKD)** from the **Select volume type** list, 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 to configure an amount of available storage:
   - Select **Amount of unused storage to configure**. 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:
    - 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, 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 open systems volumes

Use this process to create open systems volumes and to specify their attributes and properties.

1. In the navigation, under Configure Storage, select Open Systems. Under Open Systems, select Volumes — Open Systems. In Volumes — Open Systems, open the **Select Action** drop-down list and select **Create...**. Then click **Go**. The Create Volume — Select extent pool page is displayed.
2. You must select an extent pool for the target volumes. You can use the **Create new extent pool** button to create a new extent pool. After you create the extent pool and the table resets, the new extent pool is available for selection.
3. After you select the extent pool, click **Next**. The Create Volume — Define volume characteristics page is displayed.
4. Define the characteristics for the target volumes. You can select any number of volume groups from the **Select volume groups** list to associate with the target volumes.
   a. Select the volume type, volume groups, and optionally select the **Enable write cache with mirroring** selection.
   b. Optionally, use the **Create new group** button (selected by default) to create a new volume group.
5. After you define the volume characteristics, click **Next**. The Create Volume — Define volume properties page is displayed.
6. Define the volume properties. Use the **Calculate max quantity** button to populate the **Quantity** field with the calculated value. To calculate the maximum size, enter a value in the **quantity** field and click the **Calculate max size** button to see the maximum size. To calculate the maximum quantity, enter a value in the **Size** field and click the **Calculate max quantity** button to see the maximum quantity.
7. If you select the **Calculate max size** button, the **Size** field is populated with the calculated value. You can overwrite the value and enter a lesser size, and you can enter only an integer. 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.

8. 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. The **Available storage in extent pool** field displays the amount of usable storage in this extent pool in GB.
9. After you define the volume properties, click **Next**. The Create Volume — Create volume nicknames page is displayed.
10. Optionally, you can 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.
11. In Verification, review the attributes and values to verify that they are correct.
12. 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 volume creation process.

Creating open systems volume groups

Use this process to create open systems volume groups and to specify their properties and parameters.

1. In the navigation, under Configure Storage, select Open Systems. Under Open Systems, select Volume Groups. In Volume Groups — Main Page, open the **Select Action** drop-down list and select **Create...**. Then click **Go**. The Create Volume Group — Define volume group properties page is displayed.
2. Define the properties. Both the **Nickname** and **Accessed by host types** fields are required. Select the appropriate host type. All valid host types will be automatically selected.
3. After you define the properties, click **Next**. The Create Volume Group — Select host attachments page is displayed.
4. Optionally, you can select the host attachment for the target volumes. You can use the **Create new host attachment** button to create a new host attachment. After you have created the new host attachment, the table resets and the new host attachment is available for selection.
5. After selecting the host attachment, click **Next**. The Create Volume Group — Select volumes for groups page is displayed.
6. 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. You can use the **Create new volumes** button to create a new volume. After you create the volume and the table resets, the new volume is 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.

Note: If an unexpected error occurs, you must click **Cancel** to exit the wizard.

7. After selecting the volumes for the volume group, click **Next**. The Create Volume Group — Verification page is displayed.
8. Use the Verification page to review the attributes and verify that they are correct.
9. 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 volume group creation process.

Applying a configuration (Simulated only)

Use this process to select, authenticate, and apply a storage unit configuration.

1. In the navigation, under Manage hardware, select Storage units. In Storage units — Main Page, open the **Select Action** drop-down list and select **Apply configuration...**. Then click **Go**. The Apply Configuration — Select application method page is displayed.
2. 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 network connected 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 Authentication page is displayed. Click **Next** to continue and go to step 4.
3. Use the Select storage complex page to connect directly to a storage unit. Click **Next** to continue. The Authentication page is displayed.
4. In the Authentication 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.
5. In the Select storage unit page, you connect directly to a storage unit. Specify the values as appropriate and click **Next**.
6. Use the Verification page to review the attributes and verify that they are correct.
7. 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.

Getting started with Copy Services

This topic provides information about requirements and guidelines for using the point-in-time copy and remote mirror and copy features of Copy Services.

You can use the DS CLI or DS Storage Manager (GUI) to perform Copy Services tasks. An Information Center (an online help system) is provided for browsing and searching online product-related documentation. To use the information center, click the question mark (?) icon that appears in the top right corner of the DS Storage Manager.

Note: For a listing of Copy Services commands, see the Reference section of the DS6000 Information Center. For a listing of Copy Services tasks that you can perform from the DS Storage Manager, see the Managing section of the DS6000 Information Center.

The following rules apply when using Copy Services functions:

1. **One or more storage units must be assigned.** Ensure that one or more storage units are configured, assigned, and operating in a normal state. See "Storage Units — Main Page" for more information. The number of required storage units depends on the function. For example, FlashCopy operations require one storage unit, but Metro Mirror and Global Mirror require two.

   Note: If you plan to use Remote FlashCopy (known as Inband FlashCopy commands on the ESS 2105), two storage units are required for this configuration.

2. **Physical connection must be established between two storage units.** If you plan to use remote mirror and copy functions, such as Metro Mirror, Global Copy or Global Mirror), ensure that a physical connection is established between two storage units. Two (or more) storage units can be connected using a fibre channel direct connection or connect through a switch. To connect the storage units, it is recommended that you have one cable from c0 to c0 and one from c1 to c1, for example, and that you have the proper port topology configuration for those connections. To configure I/O ports, in the navigation, select **Real-time Manager** → **Storage units** → **Select Action: Configure I/O Ports...** → **Go**.

3. **Logical configuration must be created.** Consider the following requirements:

   a. **Volume capacity:** Ensure that the capacity of your target volumes is equal to or greater than your source volumes. When you select target volumes from the DS Storage Manager, it verifies that the capacities of the target volumes are at least as big as the source volumes. It does not allow you to select smaller-sized target volumes.

      Note: Be aware that for failover and failback operations to complete successfully, the volumes must be the same size and type.

   b. **Volume quantity:** Ensure that you have at least one target volume for each source volume that is of equal or greater capacity than the source volume. You can create up to 256 volumes per LSS.

   c. **Volume sizes:** Capacities of the volume are configured using the following conventions:

      Decimal
      1 GB (10^{9}) = 1,000,000,000 bytes (ESS 2105 volumes are configured in decimal format.)

      Binary
      1 GB (2^{30}) = 1,073,741,824 bytes (DS volumes are configured in binary format.)

      Block 1GB = (2^{30}) = 1,073,741,824 (iSeries™ volumes are configured in this format.)

      Note: You must consider the gigabyte definitions. In many applications, the source and target of a remote mirror and copy relationship must be exactly the same size. For example, if you plan to use DS6000 and

ESS 2105 volumes for remote mirror and copy functions, the volumes on the DS6000 must be created in decimal format to be compatible with ESS volumes.

Consider the following capacity descriptions:

Binary
This method provides volumes that fully use the capacity in every extent and also are supported across the DS6000 product families, but not the ESS 2105.

Decimal
This method provides volumes that are supported across the DS6000 and the ESS 2105 product families.

Block
This method supports volume capacity in bytes (512-byte logical blocks). Supported storage sizes range from 1 to 4G blocks (the actual number of gigabytes is the number of blocks times 512).

d. **Logical subsystem:** You can configure up to 32 LSSs. Each LSS is made up of either CKD or FB volumes. An LSS that consists of CKD addresses requires that other LSSs also be made up of CKD addresses. You can have both CKD and FB LSSs on the same storage unit.

Note: CKD LSSs are referred to as LCUs in the DS Storage Manager.

4. **Paths must be created:** You must define paths for Metro Mirror, Global Copy, and Global Mirror functions. Fibre Channel is used as the communications link between source and target volumes. To create paths, in the navigation, select **Real-time Manager** → **Copy Services** → **Paths.** From the Select Action drop-down list, select **Create...** and then **Go**. See Creating remote mirror and copy paths for more information.
5. **Relationships must be created:** Determine which source and target volumes you wish to pair for Copy Services relationships. To create relationships, in the navigation, select **Real-time Manager** → **Copy Services** → **select the function (FlashCopy, Metro Mirror, or Global Mirror)**. From the Select Action drop-down list, select **Create...** and then **Go**. See Creating FlashCopy volume pairs or Creating Metro Mirror volume pairs, for example.

z/OS® Global Mirror limitation:

If you plan to use z/OS Global Mirror (previously known as Extended Remote Copy or XRC), be aware that a z/OS Global Mirror environment that includes a DS8000 as a primary storage unit and a DS6000 as a secondary storage unit is not recommended for failover and failback operations because of the following limitations:

Performance mismatch (mirroring)
If the secondary storage unit, the DS6000, and its connectivity to the System Data Mover (SDM) that runs on z/OS Global Mirror, is significantly less capable (lower performing) than the primary storage unit and its connectivity to the application systems, the overall z/OS Global Mirror performance may suffer degraded performance. That is, if applications can write faster to primary storage units than the SDM can write to the secondary storage units, then implementations problems will result. (The SDM is the function that copies data from the primary storage unit to the secondary storage unit in a z/OS Global Mirror environment.)

Performance mismatch (running applications)
Suppose a disaster or failure occurs and applications failover to the secondary (or recovery) site and are running using the secondary storage units. If the secondary storage unit, the DS6000, is less capable (performance wise) than the primary storage unit, it is likely that you will *not* be able to complete primary business applications in the required or expected time frame.

z/OS Global Mirror-capable primary storage units
Suppose a disaster or failure occurs in an z/OS Global Mirror environment and applications failover to the secondary site and are running at the secondary site on the secondary storage units. Later, after the primary site has been repaired and is ready to resume as the primary site, the secondary storage unit can then use z/OS Global Mirror to failback to the primary site. However, for the failover and failback operations to work successfully, the secondary storage unit must be a z/OS Global Mirror-capable primary storage unit, which means it must be capable of being an z/OS Global Mirror primary storage unit. The DS6000 does not have the appropriate microcode functionality to be a z/OS Global Mirror-capable primary storage unit; and therefore, cannot be used to failback to the primary site.

General considerations include:

- If you plan to issue DS6000 commands, you must have the DS CLI prompt, and be connected to a storage unit that will be used for open systems or zSeries® host system storage. The DS CLI helps enable open systems hosts to invoke and manage FlashCopy and remote mirror and copy operations through batch processes and scripts. For more information, see the *IBM TotalStorage DS6000 Command-Line Interface Guide.*

  Note: For more complex Copy Services environments, you might find invoking and managing Copy Services functions with the DS CLI is easier. With the DS CLI, you can save commands as scripts, which significantly reduces the time to create, edit and verify their content.
- You can use Global Mirror to create consistent copies of your data at a secondary site, with minimal impact to the primary site. Global Mirror uses the concept of *sessions* to internally manage data consistency across storage units. You can also use Metro Mirror, Global Copy, and FlashCopy (without Global Mirror) to create data consistency. However, this requires that you use either external automated software or manually suspend your applications at the local site to create consistency at the secondary site.
- The DS Storage Manager can be used for almost all functions for Copy Services. However, you cannot issue the following functions from the DS Storage Manager. They are available only through the DS CLI:

  FlashCopy consistency groups
  Consistency group commands allow the storage unit to freeze I/O activity to a LUN or volume until you issue the FlashCopy consistency group command. Consistency groups help create a consistent point-in-time copy across multiple LUNs or volumes, and even across multiple storage units.

  Remote FlashCopy (known as Inband FlashCopy commands on the ESS 2105) Remote FlashCopy commands are issued to a source volume of a remote mirror and copy volume pair on a local storage unit and sent across paths (acting as a conduit) to a remote storage unit to enable a

FlashCopy pair to be established at the remote site. This eliminates the need for a network connection to the remote site solely for the management of FlashCopy.

- If you perform scenarios that call for freeze and run operations for remote mirror and copy operations, you must issue these requests from the command line interface or use automated software. These requests are *not* supported by the DS Storage Manager. (Automation software is not provided with the storage unit; it must be supplied by the user. However, IBM has offerings to assist with this automation. For more information, contact your IBM storage representative.)

Chapter 8. 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.

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

Note:

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

Preparing for the IBM TotalStorage DS CLI installation

The IBM TotalStorage 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 silent mode, console mode, or graphic mode. The mode of installation that you choose is often dependent on system configuration.

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

- Before you use the DS CLI commands, be aware of the following requirements:
  - Your Storage 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.

    For additional information, see "Creating a real-time storage configuration that includes open systems volumes" in the *IBM TotalStorage DS6000 Installation, Troubleshooting, and Recovery Guide* or this same topic in your online information center.
  - 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
Sun Solaris (7, 8, 9)	setupsolarisSparc.bin
HP Tru64 (5.1, 5.1A)	setupgenericunix.sh
Novell Netware 6.5	setupwin32.exe

Supported Host Systems	Installation File Name
iSeries OS/400 5.2	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 **Note:** Use the same installation file name for your host operating system as shown in the rows above, but add the -console parameter after the installation file name. For example: setupaix.bin -console setuplinux.bin -console	setup*os*.bin -console **Note:** Where *os* represents the name of your operating system as shown in the rows above.

- 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 1.4.1, which you can install if your system is not equipped with this version of Java.
- OS/400 installation tips:
  - Ensure that the machine that you use to install DS CLI onto a machine that is running OS/400 is network-attached to the OS/400 machine and is a supported platform for DS CLI. It is not possible to install the DS CLI program directly onto a machine that is running OS/400.
  - Provide the machine name, username, and password when prompted.
- 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.
- The installation process installs the DS CLI in the following default directory:

  AIX /opt/ibm/dscli

  HPUX /opt/ibm/dscli

  Sun Solaris
  /opt/ibm/dscli

  Windows
  C:\Program Files\IBM\dscli

  HP Tru64
  /opt/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* is 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>
   ```

   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 silent, console, or graphic installation.

Installing the IBM TotalStorage DS CLI using silent mode

The DS CLI can be installed silently (without prompts or feedback) from the command line. You must create an options/response file and use a text editor to change the default selections to a selection of your choosing. Perform these steps to install the DS CLI in silent mode.

Note: When you are installing on a Novell system, run the windows installer on a Windows system, select the Novell install option, and then select the destination where the Novell drive is mapped.

1. Log on to your system as an administrator.
2. Insert the DS CLI installation CD into the CD-ROM drive. If the Installshield starts, click the **Exit** selection to exit.

3. Generate the template response file. A response file is provided in the cliReadmes directory of the CD. You can copy it to a directory of your choosing. The following is an example of an response or options file.

   Note:

   - The response file that you want to edit is included on the DS CLI CD. Open the file with a text editor and modify the default options with the values that you want. If you are installing on a Novell system, you must indicate where the Java code is installed. The Java code must be JVM 1.4.1 or later.
   - If your system does not have Java 1.4.1 or later, the installation fails. However, Java 1.4.1 is contained on the DS CLI CD and can be installed on your system. After installation of this Java version, you can continue with the installation of the DS CLI.
   - If you are installing on an OS/400 system, the installer prompts you to provide the machine name, user name, and password for the OS/400 machine.

Example response 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"

##########################################################################
```

The following is a continuation of the response file found on the installation CD.

```
 IBM TotalStorage 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 OS/400 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>"
#
##############################################################################
```

4. Save the modifications to the response file.
5. Type the following command for all systems but the OS/400 at the command prompt and press the **Enter** key on your keyboard to start the installation process in silent mode: *setup<platform>. <exe|bin|sh> -silent -options options.txt*
6. Type the following command if you are installing onto an OS/400 system and press the **Enter** key on your keyboard to start the installation process in silent mode: *setup<platform>. <exe|bin|sh> -os400 -silent -options options.txt*

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.

Installing the IBM TotalStorage DS CLI using console mode

UNIX users who do not have an X display primarily use the console mode of installation to install the DS CLI. However, users can also run the installer from a command prompt on a Windows system.

The console mode installer displays its various screens as text. The screens have exactly the same options and information as the GUI installer. Perform the following steps to install the DS CLI using the console mode.

Note:

- The installation process checks for a version of Java 1.4.1 or higher. If this version of Java is not present the installation process stops.
- The DS CLI install CD-ROM contains the Java 1.4.1 version. You can install this version from the CD and continue with the installation of DS CLI.
- If you are ultimately installing onto an OS/400 system, ensure that the machine that you are installing from is network-attached to the OS/400 machine.
- Ensure that your Internet browser has animations turned on if you want to observe the installation progress bars that are associated with the installation of the DS CLI. Use the following instructions for your specific browser:
  - Internet Explorer
    1. From the **Tools** menu, select **Internet Options**.
    2. Select the **Advanced** tab and scroll down to the **Multimedia** section.
    3. Check **Play animation in web pages**.
  - Netscape
    1. From the **Edit** menu, select **Preferences**.
    2. Double-click on **Privacy and Security**.
    3. Select **Images** and select **as many times as the image specifies** in the **Animated image should loop** section.

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

1. Open a command prompt and locate the setupwin32console.exe file on the DS CLI CD.
2. Type the following command on the command line (unless you are installing onto an OS/400 system): setup *platform*<.exe | .bin | .sh> -console

   For example, for Windows, type: `setupwin32console.exe`

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

   Note:
   - When you issue the setupwin32console.exe command on a Windows system, you do not need to include the -console argument as part of your command.
   - If you are installing onto an OS/400 system, you are prompted to provide the OS/400 machine name, user name, and password.
3. The Welcome screen is displayed. Press `1` for Next, `3` to Cancel, or `4` to Redisplay.

```
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 TotalStorage DS Command-Line
Interface (CLI)
The InstallShield Wizard installs IBM TotalStorage 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 cliReames/ directory of this
CD-ROM.


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

4. The License Agreement screen is displayed. Press 1 to accept the terms of the license agreement, and then press 0 to notify the installer that you are finished with this section. Press 1 to continue. The following screen is displayed.

```
Please read the following license agreement carefully.

Use of the IBM TotalStorage 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]

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

5. The Windows User's screen is displayed if you are installing on a Windows system or are pointing to a Novell system. If you want Windows, press 1. If you want the Novell system, press 2. If you choose 1 or 2, the screen is updated with an X next to the system that you selected. Press 0 to notify the installer that you are finished with this section, and then press 1 to continue.

```
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. The Installation Location screen (Windows was selected) is displayed. (If you selected Novell, go to the next step) . If you are satisfied with the directory shown in brackets, press 1 to continue. If you are not satisfied, enter the directory where you want to install the CLI, and press 1 to continue. Avoid using blank spaces in the path.

```
IBM TotalStorage DS Command-Line Interface Install Location

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

Directory Name: [C:\Program Files\ibm\dscli]

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

7. If you selected Novell, the following screen is displayed.

   Note:
   - 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:
   - J: (could be any drive letter)

```
IBM TotalStorage 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]
```

8. If you selected Windows and are installing onto an OS/400, the following screen is displayed:

```
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]
```

9. The Installation Preview screen is displayed. Depending on your system's current configuration, the total size of the installation can vary. Press 1 to continue.

```
IBM TotalStorage 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]
```

10. The Installation Progress screen is displayed. Press 1 to continue after the installation completes.

```
Installing command-line Interface. Please wait...

Installing IBM TotalStorage DS Command-Line Interface. Please wait...



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


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



Creating uninstaller...
```

11. The Installation Completed screen is displayed. If you are satisfied with the installation, press `1` to continue.

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

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

12. The Readme screen displays. It contains specific information about the DS CLI for your operating system. 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 and then press `3` to finish the installation process.

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

```
To complete the installation you must restart your computer.

      [X] 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:
```

13. Select the item 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 IBM TotalStorage DS CLI using graphical mode

Users of Windows, Novell, or UNIX systems can use the DS CLI graphical mode to install the DS CLI.

You can install the IBM TotalStorage DS CLI using the graphical mode with the help of an installation wizard. Ensure that your system has the correct version of Java (Java 1.4.1) or higher and perform the following steps to install the DS CLI using the graphical mode.

Note:

- The installation process stops if the correct version of Java is not found. However, the DS CLI installation CD contains the Java 1.4.1 version. You can install this version on your system and continue with the installation of the DS CLI.
- If you are installing on a Novell system, you are directed to provide information about where your Windows drive is mapped and where the JAVA HOME location is located.
- You cannot install the DS CLI on a Windows 64-bit operating system.
- If you are installing onto an OS/400 system, ensure that the machine that you are installing from is network-attached to the OS/400 machine.
- Ensure that your Internet browser has animations turned on if you want to observe the installation progress bars that are associated with the installation of the DS CLI. Use the following instructions for your specific browser:
  - Internet Explorer
    1. From the **Tools** menu, select **Internet Options**.
    2. Select the **Advanced** tab and scroll down to the **Multimedia** section.
    3. Check **Play animation in web pages**.
  - Netscape
    1. From the **Edit** menu, select **Preferences**.
    2. Double-click on **Privacy and Security**.
    3. Select **Images** and select **as many times as the image specifies** in the **Animated image should loop** section.

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

   If you are ultimately installing onto an OS/400 system, add -os400 to your command that includes the setup file. For example, `setupwin32.exe -os400`
2. The Welcome window is displayed. Click **Next** to continue or **Cancel** to exit the installation.

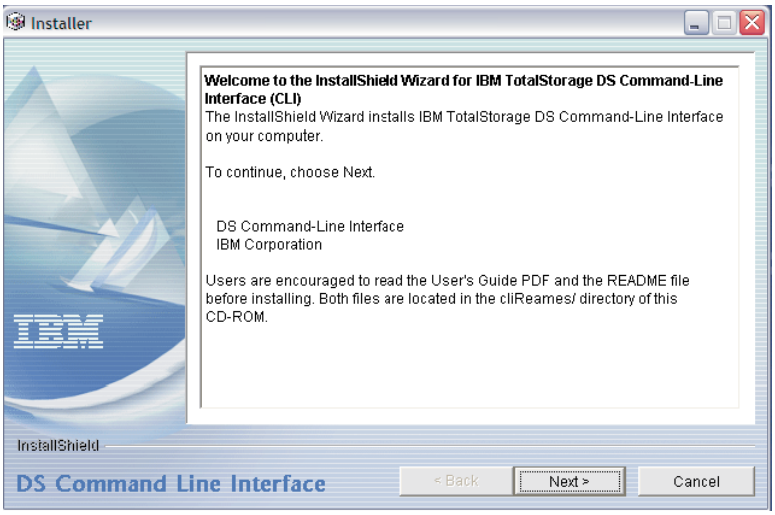

3. The License Agreement window is displayed. Select ″**I accept the terms of this license agreement**″ to continue. Click ″**I do not accept the terms of this license agreement**″ or **Cancel** to exit the installation.

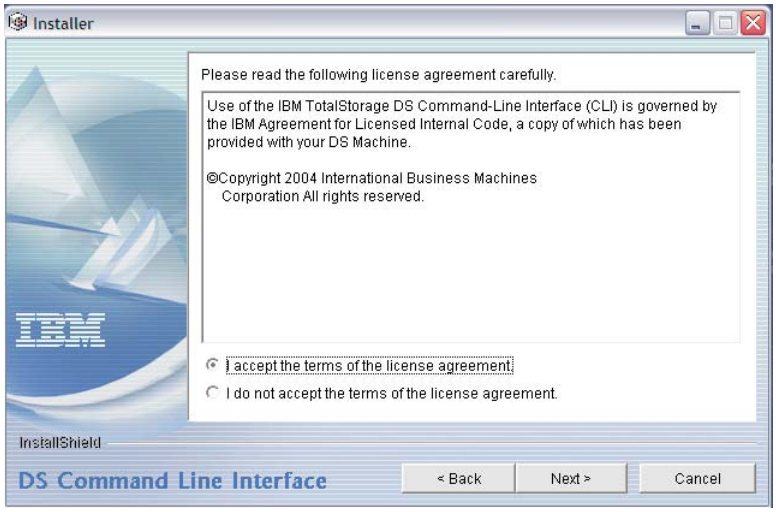

4. The User's window is displayed. Select the appropriate target system where you want the DS CLI installed, and then click **Next** to continue or **Cancel** to exit the installation.

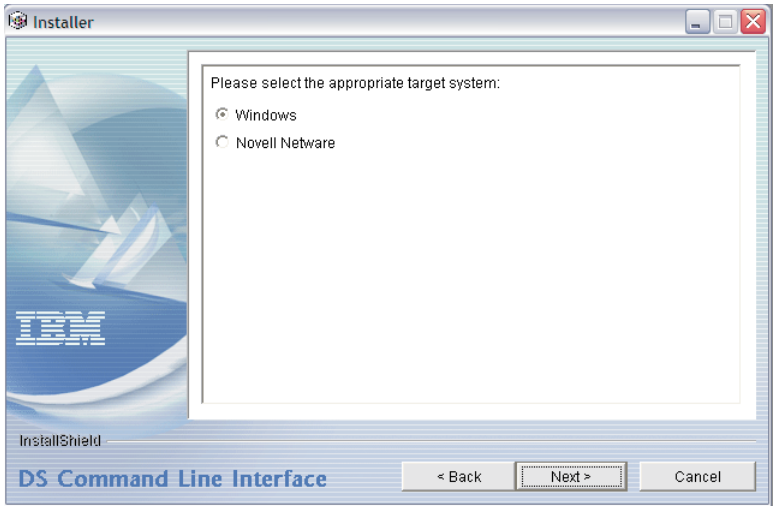

5. The Installation Location window (Windows has been selected as your system) is displayed. If you selected Novell, go to the next step. If you are satisfied with the default directory, click **Next**. If not, type the directory path where you want to install the DS CLI and click **Next**. Try to avoid using white spaces in the installation path. Click **Cancel** if you want to exit the installation.

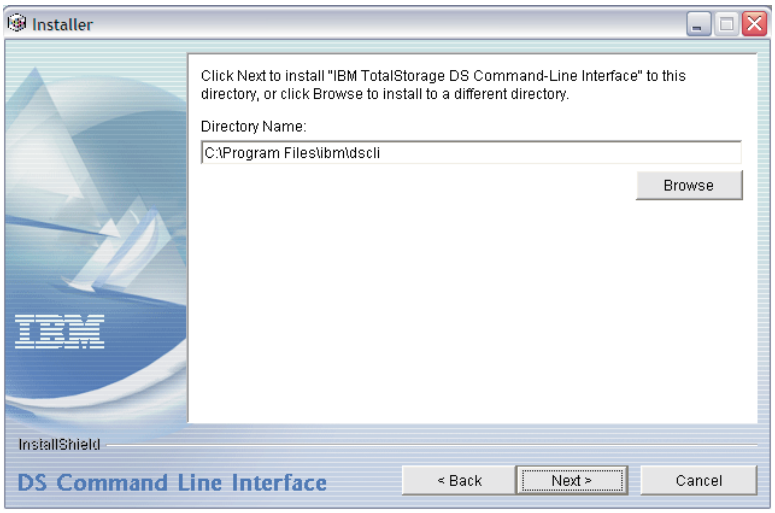

Note: If you are installing onto an OS/400 system, a window asking for the directory where Java is installed on the OS/400 is displayed when you click **Next**. Go to Step 7.

6. The Novell Location window is displayed. Complete the information for where the Windows drive is installed and where JAVA HOME is located. Click **Cancel** if you want to exit the installation.

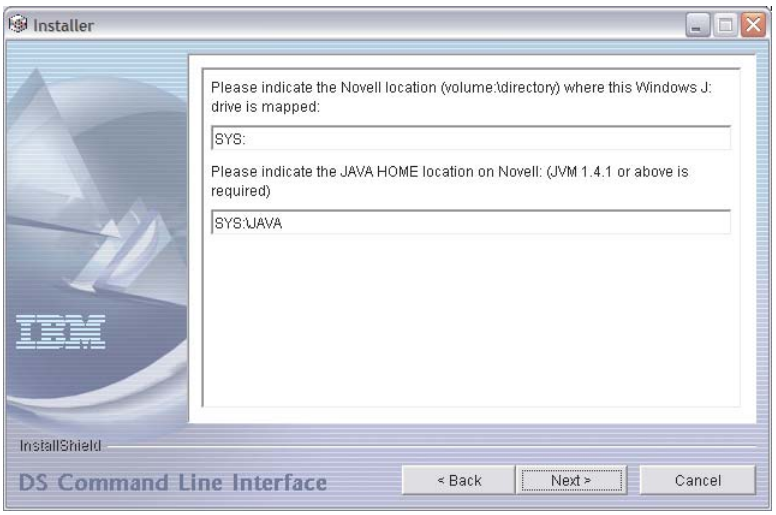

7. The Java directory window for OS/400 is displayed. Insert the directory information. Click **Next** to continue or **Cancel** to exit the installation.

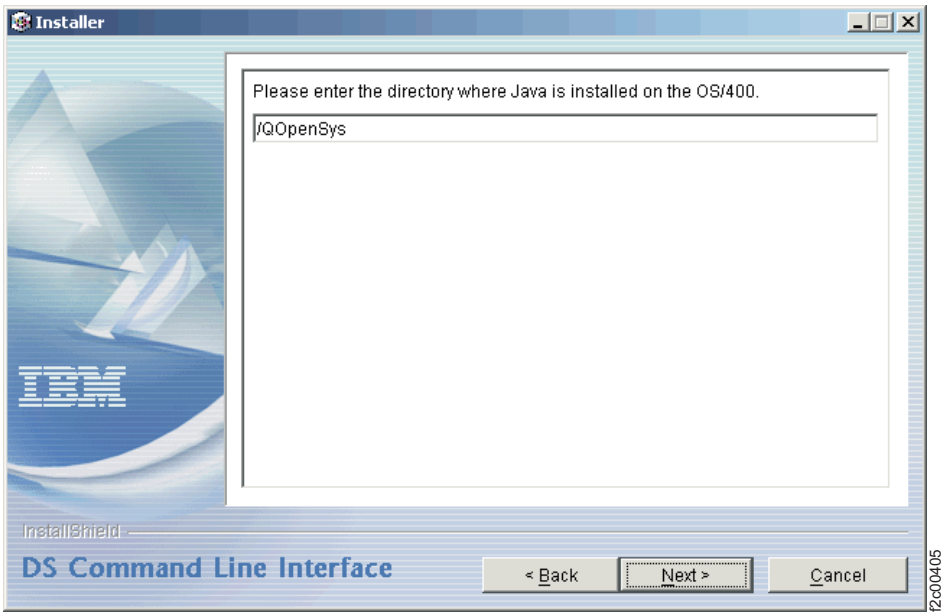

8. The Installation Preview window is displayed. This window displays where the command-line interface will be installed and how much space it will occupy on your drive. Click **Next** to continue or **Cancel** to exit the installation. You can change the installation directory by clicking **Back**.

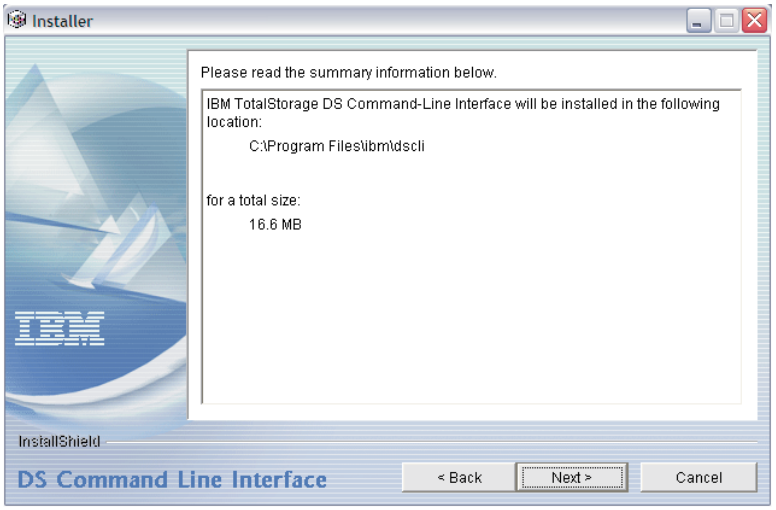

9. The Installation Progress window is displayed. This window displays the progress of the command line interface installation. Click **Next** to continue or **Cancel** to exit the installation.

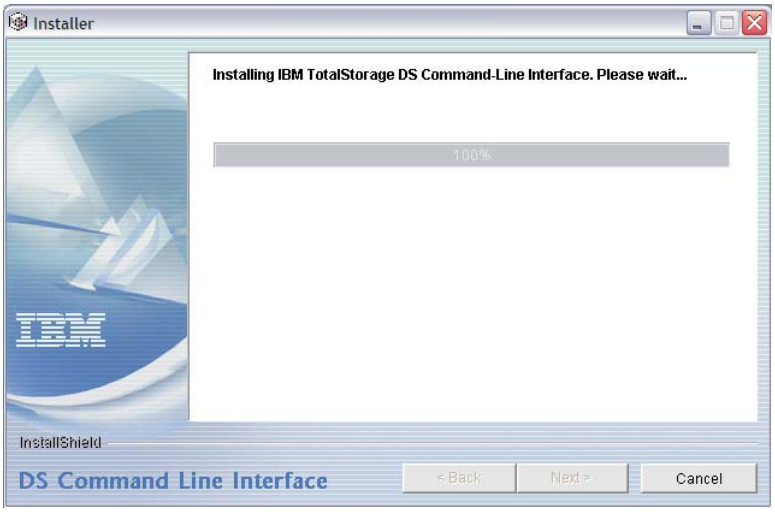

10. The Installation Summary window is displayed. This window displays the installation summary information. Click **Next** to continue or **Cancel** to exit the installation.

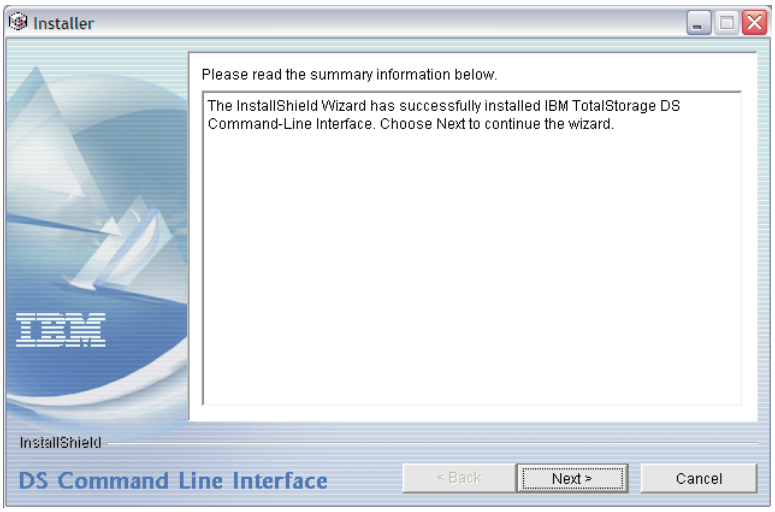

11. The Installation Complete window is displayed. A reminder is provided in text to read the Readme file. Click **Finish** to complete the installation.

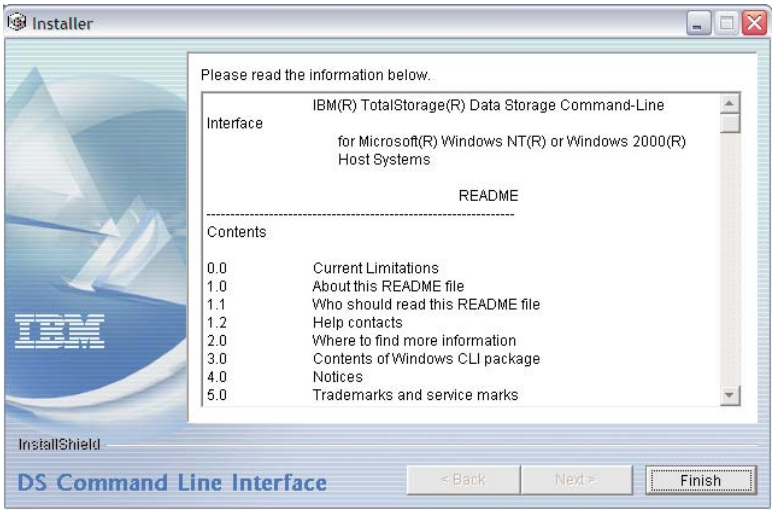

You can verify that the command line interface has installed correctly by reviewing the CLI.CFG file in the lib/ directory.

Installing DS CLI on an OpenVMS system

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.

This topic describes how to install the DS CLI on an OpenVMS system. Additionally, it provides hints for integrating the DS CLI into the operating system environment.

Preparing to install the DS CLI on an OpenVMS System

The IBM TotalStorage 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 sytem 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.
- 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.
- Ensure that you have at least 140000 blocks (approximately 70 MB) of free space on the installation disk.

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 SYSLCK, 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 DKAnnn:** where *nnn* represents the number that is assigned by the OpenVMS system to your CD device.
4. Type the following command and press `Enter` to access the command-line interface installation package in the root directory of the CD: **DIRECTORY /FULL DQA0:[000000]IBM-AXPVMS-DSCLI-*.PCSI**. 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**.

   After typing this command, 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 configuration options 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 TotalStorage DS Command-line Interface

   Copyright 2005 International Business Machines Corporation
   All rights reserved.

   Use of the IBM TotalStorage 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]
```

   Note: You can answer `Yes` if you are using the default options for your installation process.

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 instalation 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? [NO]
```

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, you will see a message similar to the following:

```
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:SYSHUTDWN.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.IBMDSCLI.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.

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%
```

4. 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. 5. Remove the command-line interface startup, login, and shutdown functions from your system startup, login, and shutdown processes.

Chapter 9. Completing DS CLI postinstallation

Perform 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 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.

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 machine and feature license codes.
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, 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, there are three command modes that are available to you.

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

- Single-shot
- Interactive
- Script

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

Sun Solaris
/opt/ibm/dscli

Windows
C:\Program Files\IBM\dscli

HP Tru64
/opt/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 is expired when 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`.
- Supply the log in information in a profile configuration file (for additional information, see the topic ″Default configuration setup with a profile file″). 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.

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. Use the following example to use the single-shot mode:

1. Enter `dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t lssi`
2. Wait for the command to process and display the end results.

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.

- 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 tucs0n -script ~/bin/mkpprcpairs`
2. Wait for the script to process and provide a report regarding the success or failure of the process.

Here is an example 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 four DS CLI commands that you have used. It also allows you to repeat any of the last four 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.

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

   Note: If you should make a mistake and type the wrong user name or password, do not try to correct this within the current session. Exit the DS CLI session 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.

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>! a list of commands is displayed 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 the third last command.

Setting up user accounts

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. 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.

Note: A user can be assigned to more than one 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.1750-75FA120
User Name testuser successfully created.
```

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, specify the output format for list commands, specify the number of rows per page in the command-line output, and specify that a banner 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 for the Windows platform .

- You can make a personal default profile by making a copy of the system default profile as <user_home>/dscli/profile/dscli.profile. The home directory, *<user_home>* is designated as follows:
  - Windows system: C:\Documents and Settings\<user_name>
- You can create a profile for the storage unit operations. Save the profile in the user profile directory. For example:
  - c:\Program Files\IBM\DSCLI\profile\operation_name1
  - c:\Program Files\IBM\DSCLI\profile\operation_name2

  These profile files can be specified using the DS CLI command parameter **-cfg *<profile_name>***. If the -cfg file is not specified, the user's default profile is used. If a user's profile does not exist, the system default profile is used.

Note: A password file, generated using the **managepwfile** command, is located at the following directory: <user_home>/dscli/security/security.dat.

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.

The profile is a text file. Table 2 provides the list of profile variables that can be used to create the profile.

Table 2. 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.
format	Specifies the output format for list commands. Specify one of the following formats: • 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. This variable is equivalent to command option -fmt. The command option -fmt overrides this default value.
fullid	Specifies that IDs display in fully qualified format, which includes the storage image ID.
header: on|off	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.

Table 2. Profile variables (continued)

Variable	Description
hmc1	Specifies the primary Storage Manager IP address. This variable is equivalent to the command option -hmc1. The command option -hmc1 overrides this default value.
hmc2	Specifies the secondary Storage Manager IP address. This variable is equivalent to the command option -hmc2. The command option -hmc2 overrides this default value.
locale	Specifies the language for the output on the local computer. • 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 2. Profile variables (continued)

Variable	Description
locale, *continued*	• 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.
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

#
# Password filename
#   The password file can be generated using mkuser command.
#
#pwfile:  ibmadmin

#
# Default target Storage Image ID
#    "devid" and "remotedevid" are equivalent to
#    "-dev storage_image_ID" and "-remotedev storeage_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  420

#
# 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

# delimitter 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
```

Activating your machine and feature licenses

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.1750-75FA120`
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.1750-75FA120
   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.1750-75FA120`
5. Press Enter and the following type of report is displayed:

   ```
   Sun Aug 11 02:23:49 PST 2004 IBM DS CLI
   ```

Activation Key	Capacity (TB)	Storage Type	Storage Facility Image
Operating Environment	45	All	IBM.1750-75FA120
FlashCopy	23	FB	IBM.1750-75FA120
...	...	...	...

Enabling remote support using the command-line interface

Remote support enables IBM support personnel to quickly assist you with problem determination and, with your consent, remotely perform certain maintenance procedures.

Enabling the DS6000 remote support function helps optimize availability of the DS6000. You can set up your site receive and send e-mail notifications as well as use the Call Home function. The Call Home function alerts IBM of a problem or a potential problem as soon as it occurs. This enables IBM Customer Service to be proactive in assisting you with maintaining your DS6000 in an optimal state. You or remote support can use the DS Storage Manager or the command-line interface to offload problem determination data to IBM.

To take advantage of remote support, you must allow an outside connection such as the following:

- A VPN connection
- An internet connection through your firewall that allows IBM to connect to your storage management system

Remote support can be set up through the DS Storage Manager or through the command-line interface. This section describes how to use the DS CLI commands to set up remote support.

Setting up contact information

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. Much, if not all, of the information that is needed for the remote support function can be collected from the configuration worksheets that are completed before your DS6000 is set up for use.

You must supply the following types of information:

- Company name
- 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 location (the building location where the parts are to be shipped)
- Ship city
- Ship state or province
- Ship postal code

- Ship country (recommended: This information is used for the Call Home e-mail records)
- Contact name (the name of the system administrator who can be contacted by IBM service personnel)
- Contact primary phone number
- Contact alternate phone number
- Contact e-mail address
- Machine identification (consists of the manufacture, machine type, and serial number)

Use the `setcontactinfo` command to provide the contact information needed for the remote support function.

1. Log into the DS CLI application for use in interactive command mode (if you are not already logged in).
2. Issue the `setcontactinfo` command from the dscli command prompt; for example, dscli> `setcontactinfo –companyname IBM –companyaddr "9000 S. Rita Rd Tucson Az 85744" –contactpriphone 1-520-7998001 IBM.1750-75FA120`
3. Press Enter to process the command. 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.1750-75FA120

   The contact information settings were successfully modified
   ```
4. Verify that the contact information is active in your system by issuing the `showcontactinfo` command as follows: dscli> `showcontactinfo IBM.1750-75FA120`

Setting up the Call Home function

The Call Home feature allows the transmission of operational and error-related data to IBM. It is the ability for the storage unit to alert IBM support to machine conditions. The information that is used for setting up the Call Home function can be collected from the worksheets that you complete before your DS6000 is set up for use.

Ensure that you have enabled the **setsmtp** command and activated the shipcountry information in the **setcontactinfo** command before you attempt to use this command.

The **setdialhome** command activates the Call Home function. You designate the machine and enable the function on the machine with this command.

1. Log in to the DS CLI application in interactive command mode, if you are not already logged in.
2. Issue the **setdialhome** command as follows: dscli> `setdialhome -action enabled IBM.1750-75FA120`
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.1750-75FA120

   The dial home settings were successfully modified
   ```
4. Verify that the Call Home function is active by issuing the **tesstcallhome** command as follows: dscli> `testcallhome IBM.1750-75FA120`
5. 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.1750-75FA120

A test problem record was successfully created.
```

Setting up call home (SMTP) notifications

One of the call home features available to you is the sending of notifications to IBM using e-mail when a problem occurs on your storage unit. You can specify that you want to enable this feature by completing the e-mail notification worksheet and by using it during the set up process.

You must specify the SMTP port and IP address where a message is sent in the event of a problem.

Use the **setsmtp** command to activate the e-mail notification remote support feature.

1. Log in to the DS CLI application in interactive command mode, if you are not already logged in.
2. Issue the **setsmtp** command as follows to activate e-mail notifications: dscli> `setsmtp –server 9.11.242.12:500 IBM.1750-75FA120`
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.1750-75FA120

The SMTP settings were successfully modified
```

Configuring new fixed block storage using the command-line interface

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.

This section describes the creation of the fixed block storage configuration and then describes the configuration of the storage unit SCSI host ports to enable access to fixed block storage. 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 ranks
- Associating ranks with extent pools
- Creating fixed block volumes
- Configuring fibre-channel I/O ports
- Creating SCSI host port connections
- Creating fixed block volume groups

Creating fixed block extent pools

Creating the fixed block extent pools is the first step in configuring new fixed block storage.

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 storage type of **fb**. The minimum requirement if for you to 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.

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. See the command descriptions for **lsextpool** and **mkextpool** if you need additional information about the commands.

1. Issue the mkextpool the following command to create the fixed block extent pool for rank group 0:

   dscli> `mkextpool -dev IBM.1750-75FA120 -rankgrp 0 -stgtype fb p000`

   Note: You must change the -dev value to match your the serial number of your system.
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.1750-75FA120

   Extent pool P1 successfully created.
   ```
3. Repeat step 1 for rank group 1, if needed.
4. Continue to repeat the mkextpool command in 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.
5. 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.

   dscli>

Creating arrays

The creation of arrays is based on the array sites that are associated with the storage unit.

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 array sites of a storage enclosure pair have identical capacity, rpm, and interface characteristics, and interface to a common DA pair.

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. See the command descriptions in Chapter 3 of the *IBM TotalStorage DS: Command-Line Interface User's Guide* if you need additional information about the **lsarraysite** and **mkarray** commands.

1. Create a list of array site IDs for all installed array sites with the status "unassigned" by issuing the following lsarraysite command:

   dscli> `lsarraysite -dev IBM.1750-75FA120 -state 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. Create an array from each site with the status ″unassigned″ by issuing the following mkarray command:

   dscli> `mkarray -dev IBM.1750–75FA120 -raidtype 10 -arsite S10`

   Repeat this command until all "unassigned" array sites have been assigned to an array.

   Note:

   - You can specify one or two array sites for Raid types 5 and 10. If there are two array sites, both 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.
   - The new array site inherits the capacity, rpm, interface, and DA pair characteristics of its parent array site. The status of the array is "unassigned" until it is assigned to a rank.

Creating fixed block volumes

To create fixed block volume groups, 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.

A storage unit contains 64,000 possible logical volumes. The quantity of logical volumes that you can create in a storage unit is dependent on the machine type and model, installed raw storage capacity, the configured storage RAID type, and the size or capacity of each volume that is created.

A logical volume consists of one or more data extents that are allocated from a single extent pool. The data type of a volume is inherited from the extent pool extent storage type (fixed block or CKD) characteristic. When a fixed block volume

is created, volume attributes are further defined as a standard storage image, OS400-protected, or OS400-unprotected volume type and volume capacity in bytes (512-byte logical blocks). These volume attributes characterize the volume to the host system that will access the volume. Each volume is assigned a volume ID, which is the volume address within the storage image 64 KB address space. Host access to a volume is enabled when the volume ID is assigned to a volume group object.

Use these steps to create fixed block volumes:

1. Create a list of fixed block extent pool IDs using the following command.

   dscli> `lsextpool -dev ID -stgtype fb`

2. Determine which extent pool IDs will be used as the source for the fixed block logical volumes to be created. 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) affinity, which dictates that the LSS ID component of the volume ID must be an even or an odd number.
3. Create a list of unassigned address groups using the following command.

   dscli> `lsaddressgrp -dev ID`

   Take the following notes into consideration:

   - Determine the defined address group objects (0 - F). If the returned list is empty, then all address group objects are available to be defined (0 - 3).
   - Address groups with storage type CKD are not available for the definition of fixed block volumes. Address groups with storage type ″FB″ and number of defined LSSs less than 16 are available for new fixed block volume ID definition.
   - If an undefined address group will be used when creating new fixed block volumes, select the lowest numbered address group that is not defined.
   - 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. Create 256 fixed block volumes for the specified LSS using the following command. Specify a volume ID that has not been previously defined as a fixed block or CKD volume. Repeat this step for all of the required logical volumes for each LSS.

   dscli> `mkfbvol -dev ID -extpool ID -cap 8.6 -name DS_1_vol#h -volgrp volume_group_ID 1000-10FF`

   Take the following notes into consideration:

   - All volumes will have the same type and capacity attributes.
   - -extpool identifies a fixed block extent pool containing available data extents.
   - -cap (capacity) is 8.6 GB.
   - -name assigns a unique name to the volume. The volume name parameter can include a wild card (#d or #h) that inserts a decimal/hexadecimal volume ID value into a volume name, thereby making volume names unique when multiple volumes are created by one command execution.
   - -volgrp (**Optional**) The use of this parameter is optional at this point in the configuration process. The use of this parameter requires that a volume group ID already exist. It is possible that you have not yet created your volume group IDs. However, if you have created the IDs, you can use the -volgrp parameter to assign the new fixed block volumes to the existing volume group IDs.

- Volume ID 1000-10FF 256 volumes, starting at a fixed address group (1), LSS ID (10), and volume number (00). A volume ID is four hexadecimal characters in the format ″llxx″, where *ll* specifies an LSS ID (00 - FE) and *xx* specifies an LSS volume number (00 - FF). A volume ID must be compatible with the extent pool ID that is the specified source for volume data extents. That is, if an extent pool ID is associated with rank group 0/1, then the volume LCU ID must be an even/odd numbered value, respectively.

Creating fixed block volume groups

To create fixed block volume groups, you must have the command-line interface prompt, and you must be connected to a storage image that will be used for open systems host system storage.

A volume group identifies the set of fixed block logical volumes that are accessible by one or more SCSI host system ports. If you make volume groups prior to making fixed block volumes, then the mkfbvol command -volgrp parameter can be used to assign new fixed block volumes to a target volume group ID. Otherwise, if you make fixed block volumes prior to making volume groups, then use the mkvolgrp command and the -volume parameter to assign volumes to a new volume group ID, or use the chvolgrp command and the -volume parameter to assign volumes to an existing volume group ID.

SCSI host system access is constrained to the “identified” access mode. For “identified” access mode, a volume group ID is assigned to a SCSI host port. Only those SCSI host ports that are “registered” to a volume group ID are allowed to access the set of logical volumes 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 K volumes. The volume group type must be selected according to the addressing capability of the SCSI host system that will use the volume group.

Create and view fixed block volume groups using the following steps:

1. Create a fixed block volume group using the following command. Repeat this command for each volume group that you want to create..

   dscli> `mkvolgrp -dev` *`ID`* `-volume` *`ID1,ID2,... volume_group_name`*

   - Use the -type scsi256 flag if the volume group is constrained to 256 volume IDs. Otherwise, a volume group may contain up to 64 K volumes. Use the -type 0s400mask flag if the volume group is constrained to fixed block volume types OS400-protected or OS400-unprotected. Otherwise, the volume group is constrained to the fixed block volume type 2107.
   - The volume group name should describe the SCSI host system ports that will access volumes using this volume group ID. A SCSI host port can access one and only one volume group ID. Multiple SCSI host ports can register to access the same volume group ID. In this case, all host ports have shared access to the set of volumes contained by the volume group ID.
2. Create a list of the assigned volume group IDs using the following command.

   dscli> `lsvolgrp -dev` *`ID`*

   If the SCSI host ports have not yet been configured, save the volume group ID for use when you create the SCSI host port. If the fixed block volumes have not yet been defined, save the volume group ID for use when you create fixed block

volumes. If the SCSI host ports were configured prior to the volume groups, use the chhostconnect command to modify each SCSI host port that is configured for SCSI host system access. Add the volume group ID to enable host system access to storage image fixed block volumes.

Configuring fibre-channel I/O ports

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.

The storage image supports the fibre-channel host bus adapter (HBA) card type. For machine type 2107, HBA cards are installed in I/O enclosures, each containing up to four HBA cards. For machine type 1750, one or two HBA cards are installed in each of the two CEC assemblies. Each fibre-channel HBA card contains four I/O ports. The storage image microcode automatically creates one I/O port object to represent each HBA card I/O port. The default fibre-channel I/O port object settings enable SCSI-FCP "identified" access to fixed block volumes. You might have to modify the I/O port object settings to enable SCSI FC-AL access to FB volumes.

Use these steps to configure fibre-channel I/O ports:

1. Create a list of fibre-channel port IDs using the following command:

   dscli> `lsioport -dev ID -type fc`
2. Identify which I/O port IDs you want to access the fixed block volumes. Configure a minimum of four storage image I/O ports for SCSI host I/O operations. Select ports with physical locations on different host bus adapter (HBA) cards. If possible, locate the HBA cards in different I/O enclosures. Identify the port IDs that you want to configure for SCSI host access.
3. Set the specified I/O ports to enable the FCAL topology using the following command. I/O ports are automatically set offline and returned to the online state after configuration changes are applied

   dscli> `setioport -dev ID -topology fc-al port_ID1 ID2...IDn`

Creating SCSI host port connections

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.

The 1750 supports the "identified" access mode for SCSI host attachments, which requires that all SCSI host ports be identified to a storage unit. This is accomplished by creating a SCSI host port object for each SCSI host port that will access storage unit fixed block volumes. A SCSI host port object contains attributes that identify the SCSI host system type, the port profile, the port WWPN, the volume group ID that the port will access, and an array of storage unit I/O port IDs that the host port will log into for volume access, or an attribute to indicate that all I/O ports can be used for volume access. Object attributes also include a host port description and port nickname. Create one SCSI host port for each port that will access storage unit volumes.

Create SCSI host ports using the following command:

dscli> `mkhostconnect -dev ID -wwname wwpn -profile # -volgrp ID -ioport ID,ID,ID,ID host_name`

The command specifies the worldwide port name, the fibre-channel communication profile, the volume group ID that this host port is allowed to access, the storage image I/O ports that this port can log in to, and your name for the SCSI host system.

Note:

- You can obtain a profile value from the lsportprof command for the host type.
- You do not need to list the volume group.
- Use the chhostconnect command to map volumes to volume groups

Modifying fixed block volume groups

This scenario describes how to modify fixed block storage within a storage unit.

To modify fixed block volume groups, 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.

Adding volumes to a volume group and removing volumes from a volume group are typical storage management tasks. The volumes that are added to a volume group can be “unassigned” to a volume group, or they can be volumes that are assigned to a volume group but you want to move them to a different volume group. In either case, you are responsible for managing how the volumes are allocated to volume groups and how the volumes are reserved for future allocation. It is better that you maintain “unassigned” volumes in a volume group that is not accessible by any host system, thereby controlling the accessibility of volumes that are reserved for future allocation.

You can assign a fixed block volume to multiple volume groups. This might be necessary for some host system applications. However, damage to volume data can occur if a volume is accessed by different host systems using different file management systems. To assign a fixed block volume to multiple volume groups, perform the following steps:

1. Find the fixed block volumes that are to be assigned to a volume group using the following command.

   dscli> `lsfbvol -dev` *`ID`* `-type` *`2107`* | *`os400-protected`* | *`os400-unprotected`* `-extpool` *`ID`*

   The command creates a list of all volumes of the specified volume type within the specified extent pool. It includes only the volumes that are contained by the specified storage image.
2. Retrieve the current volume group volume map using the following command.

   dscli> `showvolgrp -dev` *`ID volume_group_ID`*

   The command creates a list of volumes that are assigned to the target volume group.
3. Modify the volume group using the following command.

   dscli> `chvolgrp -dev` *`ID`* `-action add` | `remove` | `replace -volume` *`ID,ID,...,ID volume_group_ID`*

   You can add or remove volume IDs to the list in order to add or remove volumes. This command applies the updated volume ID list.

Chapter 10. Registering for My Support

Use this process to register for My Support.

My Support provides pro-active notification of microcode updates though an email address that you specify. My Support will automatically notify you of the latest microcode fixes and how to obtain them. It is highly recommended that you register for My Support.

To access online technical support, visit: `http://www-1.ibm.com/servers/storage/support/disk/ds6800/`. My Support registration provides e-mail notifications when new firmware levels have been updated and are available for download and installation. To register for My Support:

1. Visit the following Web site: `http://www.ibm.com/support/mySupport`
2. Select `register now`.
3. Under `My IBM Registration Step 1 of 2`, fill in the required information. Items with an asterisk (*) are required fields. Select `Submit`.
4. Under `My IBM Registration Step 2 of 2`, fill in the required information. Items with an asterisk (*) are required fields. Select `Submit`. Select `Continue`.
5. Login. Enter your User ID and Password and select `Submit`.
6. Select the `Edit my profile` tab. Select the information required for your profile under the `Products` section.
   a. In the drop-down list, select `Storage`.
   b. In the drop-down list, select `Computer Storage`.
   c. In the drop-down list, select `Disk Storage Systems`.
   d. In the drop-down list, select `TotalStorage DS6000 series`.
   e. Check the check box next to `TotalStorage DS6800`.
   f. Select `Add products`.
   g. Review your profile for correctness.
7. Select the `Subscribe to email` tab. In the drop-down list, select `Storage`. Then check `Please send these documents by weekly email` and check `Downloads and drivers` and `Flashes`.
8. Under the Welcome box, select `Sign out` to end your session.

Installing a modem on the management console

Perform the following steps to install a modem on the management console. You can use this modem for VPN connections during remote support situations.

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

You can use this 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.

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. Complete 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. Use the serial cable to connect the modem to the management console. The 9-pin serial connector connects to the management console while the 25-pin serial connector connects to the modem.
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**, then **Settings**, then **Control Panel**, and finally **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 as appropriate before you 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 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. Click **Browse...** on the Install From Disk panel and navigate to the CD-ROM directory. Click **OK**.
   i. Select **MT5600BA V92** from the Models list.
   j. Select the communications port where you installed the modem.
   k. 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** drop-down 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 from the drop-down list 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.

Enabling remote support function

Use this section to enable remote support function.

Enabling the DS6000 remote support function helps optimize availability of the DS6000. Remote support enables highly-skilled IBM support personnel to quickly assist you with problem determination and, with your consent, remotely perform certain maintenance procedures. An additional function of remote support is Call Home, which alerts IBM of a problem or a potential problem as soon as it occurs. This enables IBM Customer Service to be proactive in assisting you with maintaining your DS6000 in an optimal state. You or remote support can use the DS Storage Manager or the command-line interface to offload problem determination data to IBM.

Setting up call home

Use this process to configure notifications.

You must define the customer contact information before configuring notifications.

This task enables you to define Call Home (SMTP), SNMP, and SIMs (service information messages) notifications for a storage unit. See Notification methods for detailed information about these functions.

1. Under **Manage hardware**, select **Storage units**. In Storage units - Main Page, select **Configure notifications** in the **Select Action** drop-down field. Then click **Go**. The Configure notifications - Define Call Home page is displayed.
2. Ensure that **Enable Call Home** is selected to activate Call Home. (This is checked by default.)
3. Complete the SMTP information.
   a. Enter the SMTP server host name.
   b. Enter the SMTP IP address.
   c. Enter the SMTP server ports.
   d. Click **Apply**.
4. Click **Test Call Home connection** to send a connection test and generate a problem log entry. A confirmation message is displayed.
5. Click the **SNMP** tab. The Define SNMP connection page is displayed.
6. Select Enable SNMP notification to define the SNMP connection properties for the selected storage units.
7. Specify either an IP address, a Host name, or both under SNMP trap destination.

8. Specify an SNMP community name of up to 32 characters. This field is used to authenticate requests. 'Public' is selected by default.
9. Optionally specify an SNMP system contact name of up to 32 characters. Enter a destination port. Click **Apply**.
10. Click the **zSeries** tab. The Define Service Information Messages for zSeries page is displayed.
11. Optionally select a SIM severity level in the Severity reporting level for DASD Service Information Messages field.
12. Optionally select a Media Service Information severity level in the Severity reporting level for Media Service Information Messages field.
13. Optionally select a Service Information severity level in the Severity reporting level for a Service Information Messages field.
14. Click Apply.

Initiating a remote support connection

Use this process to start remote support on the systems management console.

Follow these steps to initiate remote support:

1. From the SMC, open a DOS window. Depending on your windows version, this can be done in one of several ways:
   a. Start->Programs->Command Prompt
   b. Start->Run->cmd.exe
2. Gather the following three pieces of information:
   a. The 7 digit alpha numeric serial number of the 1750 (including plant of mfg)
   b. The IP Address of Processor Card 1
   c. The IP Address of Processor Card 2
3. Issue the following command in the command prompt window: ibmremote -s *serialnum -ip1 -ip2*, where *serialnum* is the serial number of the 1750 from step 2; ip1 is the IP address of the processor cord 1, and ip2 is the IP address of the processor cord 2.
4. Within one minute, the VPN connection is established, indicated by an IBMVPN network connection icon in the system tray. If the IBMVPN connection is not established, contact your next level of IBM support.
5. Communicate the 7 digit serial number to your next level of support.
6. Perform the following steps to stop remote support on the SMC (Systems Management Console)
   a. Right click on the IBMVPN network connection icon in the system tray.
   b. Select **Disconnect**.
   c. Verify that the IBMVPN network connection icon goes away.

Chapter 11. Removing the DS6000 Storage Manager and CLI

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.

The following topics describe the steps required to successfully remove the interfaces.

DS6000 Storage Manager

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

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

1. Navigate to your Control Panel and open the Add/Remove program facility.
2. Scroll the list of currently installed programs and click the listing for IBM TotalStorage DS6000 Storage Manager.
3. Click the **Change/Remove** button, and the Welcome window for the Uninstaller is displayed.

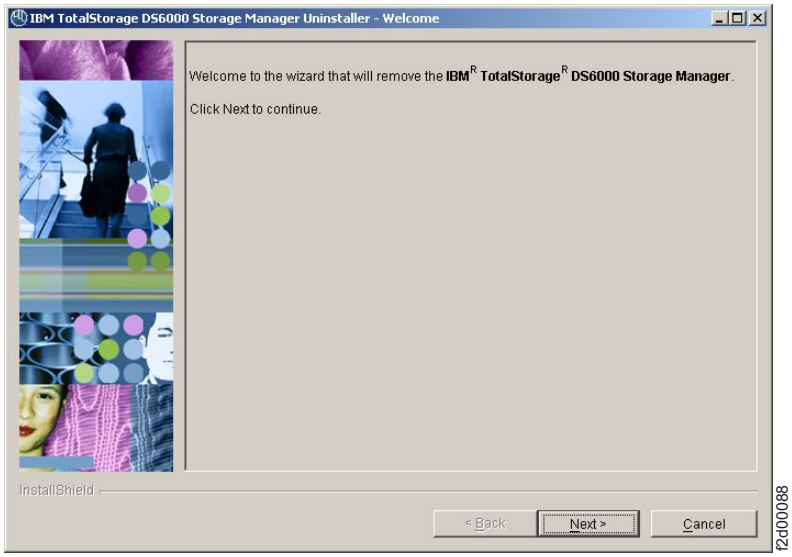

Uninstaller Welcome window

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

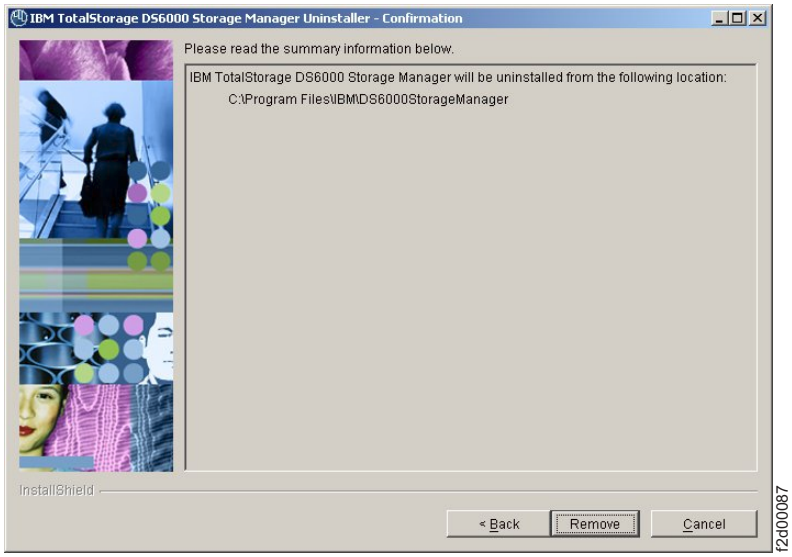

Uninstaller Windows Confirmation window

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

   During the uninstallation process, the following system checks occur:

   - A check of your system to detect if DS Storage Manager Server is used by other programs on the system and if not used, removes it from the system. Otherwise, the DS Storage Manager Server is not removed from your system.
   - A check of your system to detect if DS Network Interface Server is used by other programs on the system and if not used, removes it from the system. Otherwise, the DS Network Interface Server is not removed from your system.
   - A check to determine whether any of the files to be removed are locked because they are in use by another process. If a locked file has been detected, you can perform the following:

     a. From Windows services, set the **IBM WebSphere Application Server V5 - DS Storage Manager** and **IBM DS Network Interface Server** services to **Manual** startup type.

     b. Reboot the computer.

     c. Proceed once again with the removal of the DS Storage Manager.

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

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

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

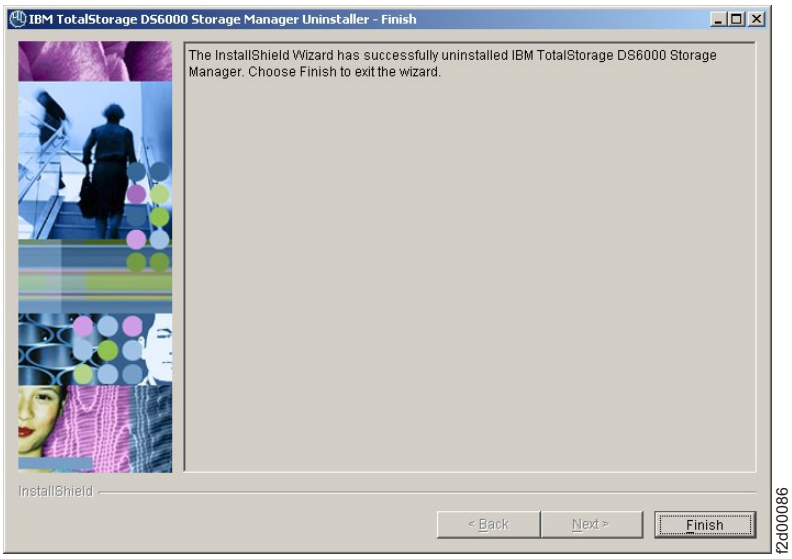

Uninstaller Finish window

6. Click **Finish** to complete the removal process and exit the GUI part of the uninstallation program.
7. Close the Add/Remove Programs window.
8. Restart your system, if required (now or later) to complete the removal process.

If you want to remove all configuration settings before reinstalling the DS Storage Manager, you can rename the IBM folder under Program Files.

Removing the DS6000 Storage Manager from a Windows operating system using the silent mode

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

1. Log on to your Windows operating system with administrator authority.
2. Open the Command Prompt window.
3. Navigate to the directory where the DS6000 Storage Manager is installed.
4. Type `<dest_path>\_uninst\uninstall.exe -silent`, where *<dest_path>* is the path where the DS6000 Storage Manager is installed.
5. Wait for the program to remove the DS6000 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 reinstalling the DS Storage Manager, you can rename the IBM folder under Program Files.

DS6000 command line interface

Removing the IBM TotalStorage DS CLI using silent mode

Use the 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, HPUX, Sun, or AIX).

Perform the following steps to successfully uninstall the DS CLI.

Note:

- If you are using Windows or Novell you will use the Add/Remove Programs feature to uninstall the DS CLI.
- 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*, with *xxx* depending 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 IBM TotalStorage 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 the 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 TotalStorage DS Command-Line
Interface (CLI)
The InstallShield Wizard uninstalls IBM TotalStorage 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 TotalStorage DS Command-Line Interface will be uninstalled from the
following location:

C:\Program Files\ibm\dscli

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 TotalStorage DS Command-Line Interface...
```

5. The Uninstallation Finish screen is displayed. Press `3` to finish the removal.

```
The InstallShield Wizard has successfully uninstalled IBM TotalStorage DS
Command-Line Interface. Choose Finish to exit the wizard.

Press 3 to Finish or 4 to Redisplay [3]
```

Removing IBM TotalStorage DS CLI from your system using graphical mode

Use the graphical mode to remove the DS CLI from your system when DS CLI is installed on a Windows, Novell, or UNIX system.

Note: The following procedure applies only to the removal of the DS CLI application. This procedure cannot be used to remove other versions of CLI.

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, you need to restart your system to complete the uninstall. Perform the following steps to remove the DS CLI using the graphical mode.

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

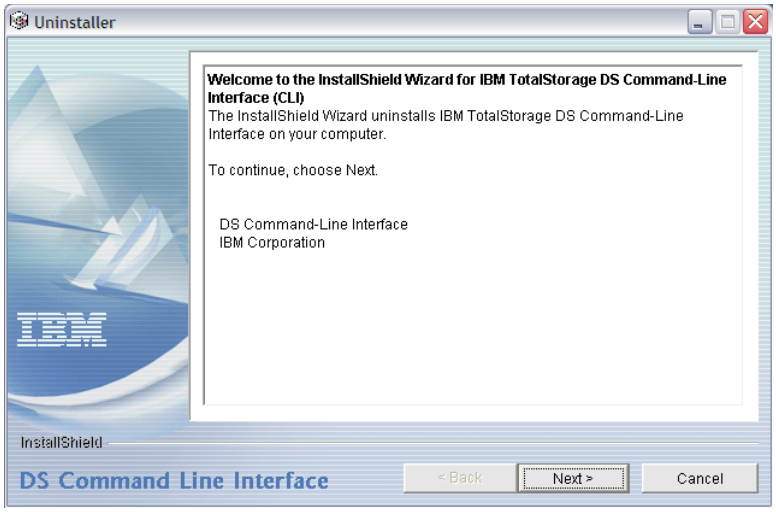

4. Click **Next** to continue or click **Cancel** to exit the removal process. When you click **Next**, the Confirmation window is displayed showing the directory from which the DS CLI program is removed.

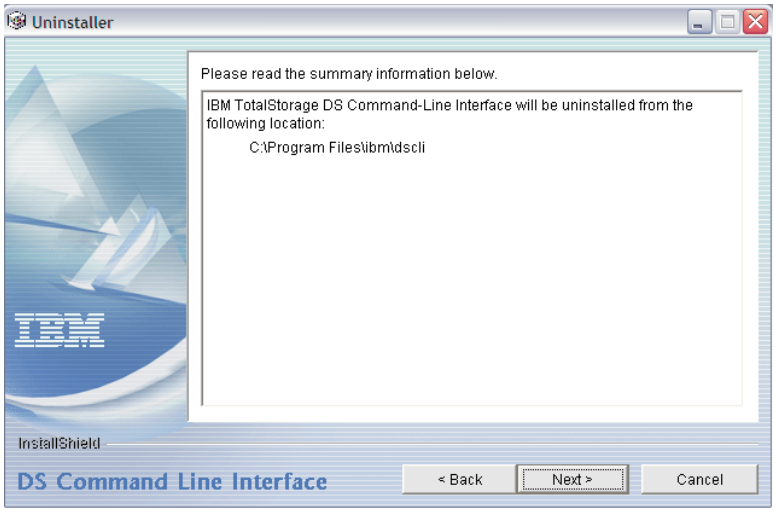

5. Click **Remove** to continue or **Cancel** to stop the removal and exit the uninstall process. Click **Back** to return to the previous window. When you click **Remove**, the Uninstallation Progress window is displayed. When the uninstall process is finished, the Finish window is displayed, which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

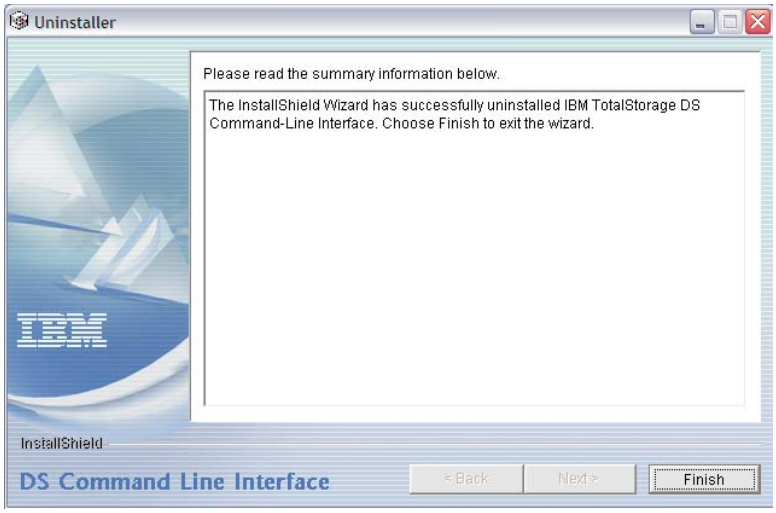

If the uninstall program does not remove some information from your system, the Restart window is displayed. You must restart so that previously locked files are released and automatically deleted.

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

Chapter 12. Upgrading the interfaces

Upgrading any of the DS6000 interfaces requires planning.

Generally, the upgrade of an interface requires that you remove the current version of the interface and reinstall the newest version of the interface using the remove and install procedures.

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

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

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

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

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

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

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

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

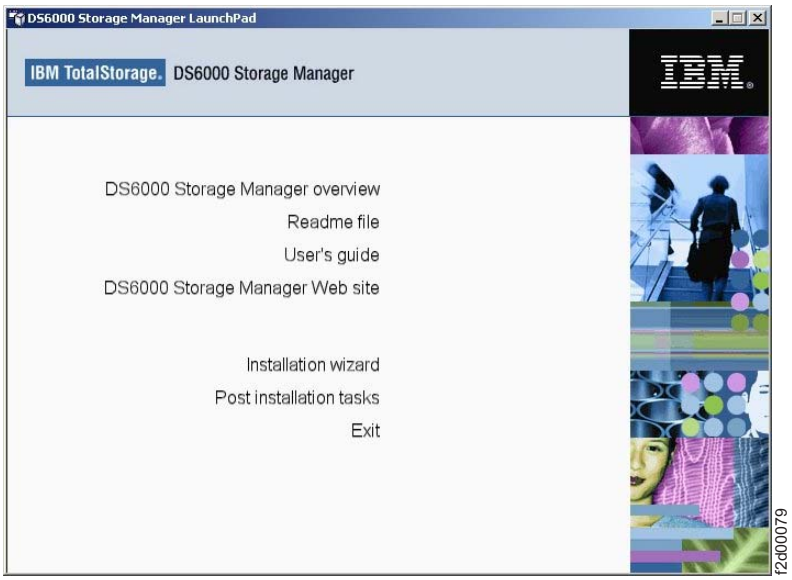

Launchpad window

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

   DS6000 Storage Manager overview
   Provides information about the IBM TotalStorage DS Storage Manager software.

   Readme file (recommended selection)
   Provides last minute product information that was not provided in these installation instructions.

   User's guide
   Provides specific installation instructions.

   DS6000 Storage Manager Web site
   Provides information from the product Web site.

   Installation wizard
   Starts the IBM TotalStorage DS Storage Manager installation program.

   Post installation tasks
   Provides information about configuring the IBM TotalStorage DS Storage Manager.

   Exit
   Exits the IBM TotalStorage DS Storage Manager LaunchPad program.

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

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

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

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

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

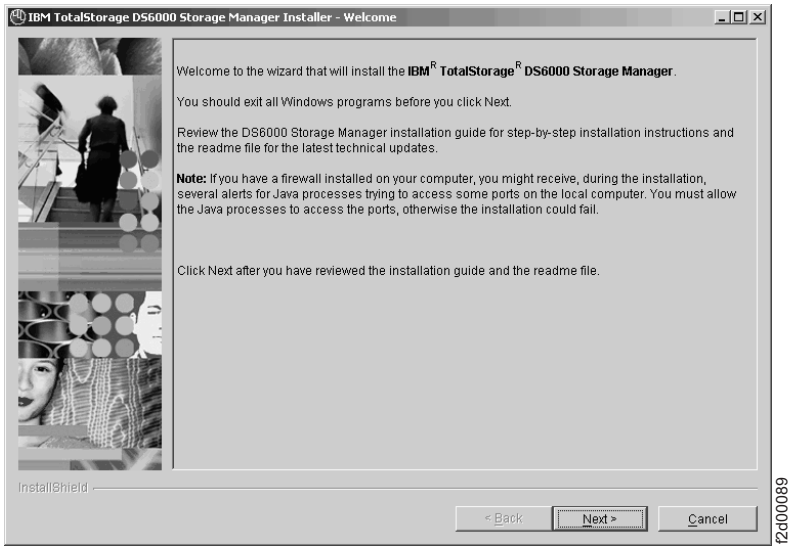

DS6000 Storage Manager Installer Welcome Window

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

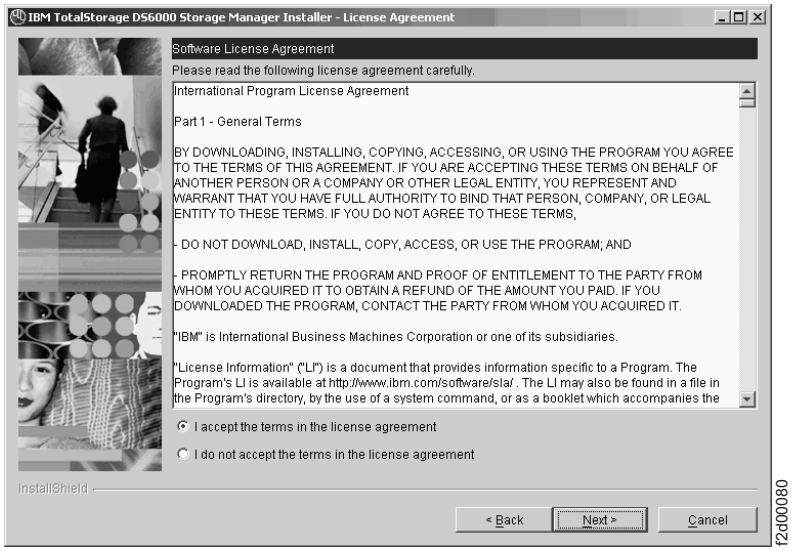

License Agreement window

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

7. Read the license agreement and click your acceptance or nonacceptance of the agreement. If you accept, the **Next** button is highlighted. Click **Next** to

continue or click **Cancel** to exit the upgrade process. When you click **Next**, the Product Version check window is displayed.

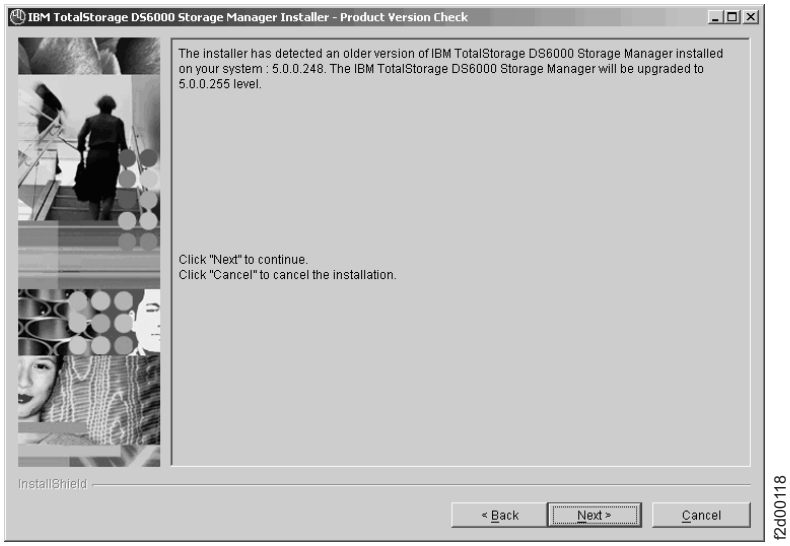

Product Version Check window

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

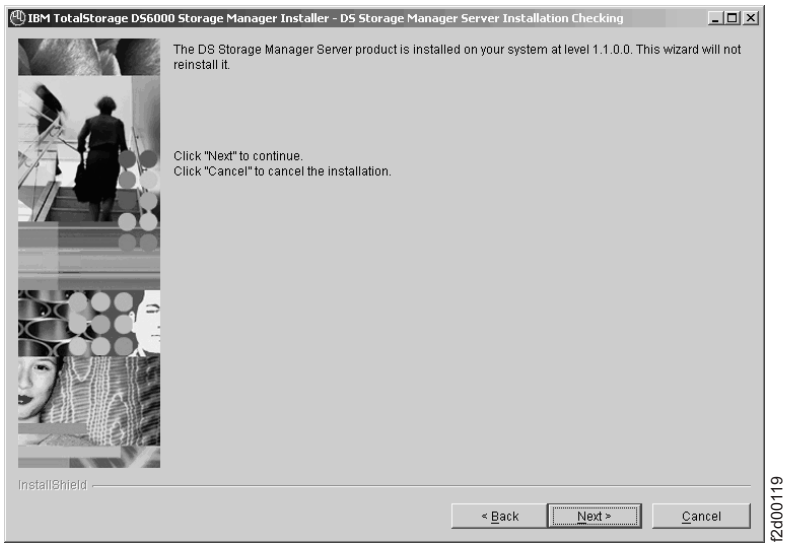

DS Storage Manager Server Installation Checking window

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

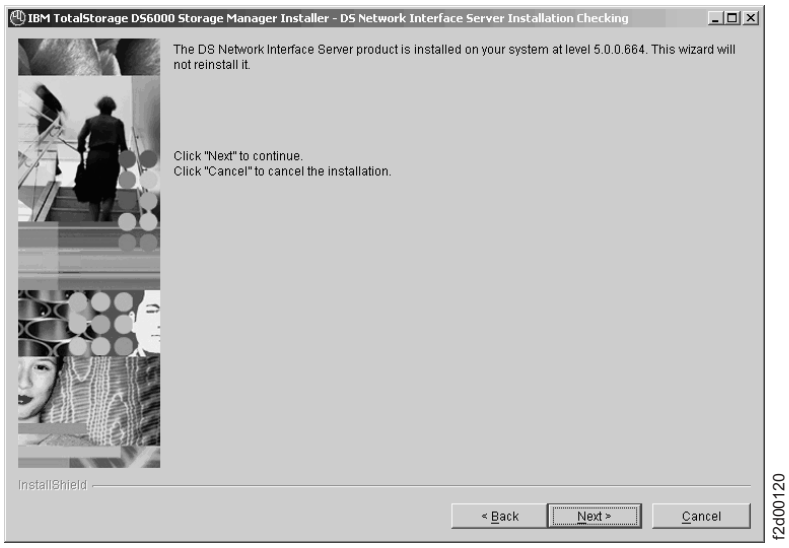

DS Network Interface Server Installation Checking window

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

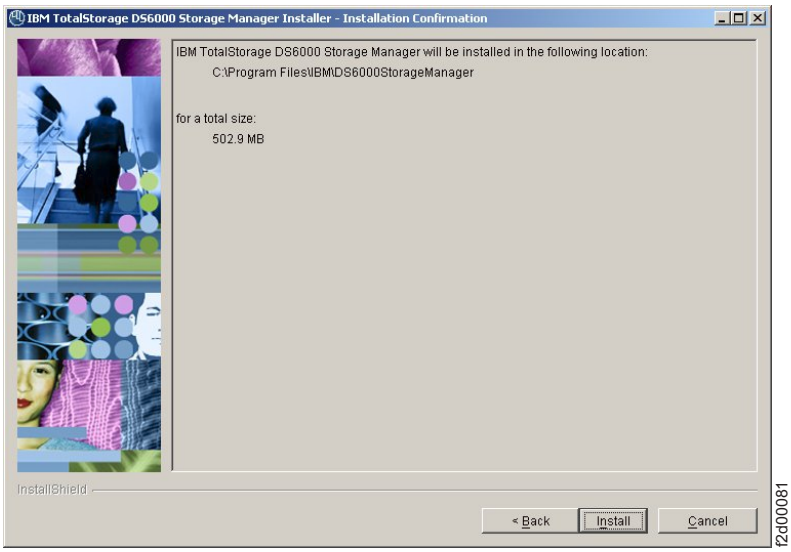

Windows Installation Confirmation window

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

    The installation process performs the following actions:

    a. If the two servers (DS Storage Manager Server and DS Network Interface Server) are already installed on your system, they are stopped in the following order of windows:

1) The Embedded IBM WebSphere Application Server - Express server (part of the DS Storage Manager Server) window is stopped first.
2) The service window (DS Network Interface Server) is stopped next.
3) The WS Help System (part of the DS Storage Manager Server) window is stopped, if it was not stopped before by the Embedded IBM WebSphere Application Server.

b. If one or both of the servers are not installed or have to be upgraded on your system, they are installed or upgraded in the following order (the progress of the installation is indicated on the associated progress window):
   1) DS Storage Manager Server Installation Progress window
   2) DS Network Interface Server Installation Progress window

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

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

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

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

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

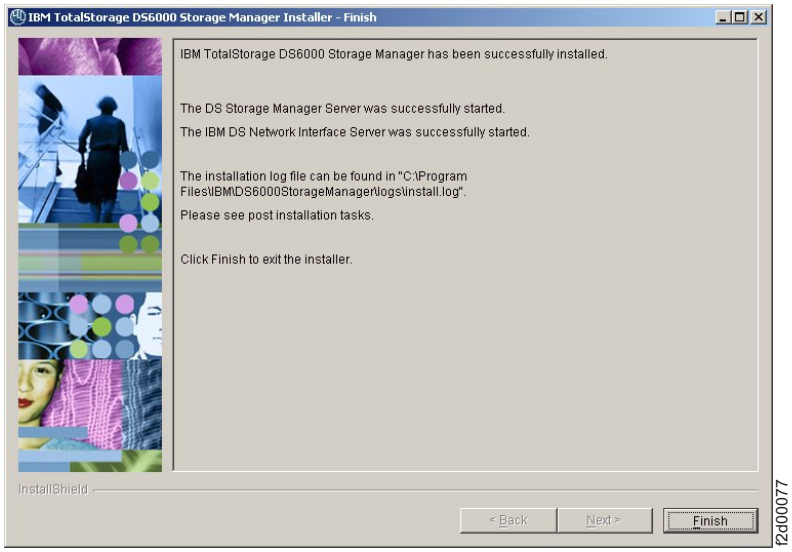

Windows DS6000 Storage Manager Installer Finish window

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

13. If the Finish window indicates that the installation fails, check the installation log for error messages. The installation log is located in *xxx\logs\install.log*, where *xxx* is the destination directory where the IBM TotalStorage DS Storage Manager is installed (for example, c:\Program Files\IBM\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 silent mode of installation.
    - If the error is due to a wrong installation parameter value, restart the installation using the steps that are described in this procedure or the steps that are described in the silent mode installation process. Navigate to the Server Parameters window and insert the correct values. Finish the installation process.
14. Complete the postinstallation tasks. If, when exiting the installation wizard, the postinstallation.txt file does not automatically open, manually open it from the LaunchPad window, and follow the instructions to complete the postinstallation tasks.
15. Exit the LaunchPad program by clicking **Exit** on the LaunchPad window. Reboot if instructed to do so.
16. You have now completed the software upgrade portion of the DS6000. Continue to the postinstallation tasks.

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

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

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

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

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

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

Upgrading your system to use DS CLI

You can upgrade your system to use DS CLI by following the remove and install procedures put in place for DS CLI. However, there are some things that you need to consider before making this upgrade.

The DS CLI is designed to support the following features that exist on the IBM TotalStorage Enterprise Storage Server® (ESS) Models 750 and 800:

- A Copy Services domain, configured as part of the IBM TotalStorage Enterprise Storage Server (ESS) Models 750 and 800.
- FlashCopy Version 2 and PPRC Version 2 licenses that are available on the ESS.

However, the DS CLI is not designed to support the CLI scripts that you have written for these features without modification.

As part of your upgrade preparation, ask yourself the following questions:

- Do you plan to add the DS6000 to your network?

  If you do, you must install the DS CLI application. This application is used for the configuration, performance, and Copy Services functions. Also, only the DS CLI can be used to perform Copy Services functions between the ESS and the DS models.
- Do you plan to keep the DS6000 separate from your network?

  If you do, you can not communicate with the ESSs in your network, even for Copy Services functions. You can install the DS CLI on your server and it supports the DS6000, but new scripts must be written that support the DS6000.

Chapter 13. 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 1GGA*n*. In contrast, fibre-channel disk devices have the OpenVMS device type DG, which result in device names in the form 1DGA*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 1GGA*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 1GGA*n* CCL devices are shown in the command output. Otherwise, the 1GGA*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 1GGA*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

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/index.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

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#processquotas.

To check if your current process quotas fulfill the recommendations, you can run the following process: `IBMDSCLI$JRE:[LIB]Java$Check_Environment.com`.

Chapter 14. 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. After going beyond the set number of allowable attempts with the wrong password, the account is locked. If the Admin account is involved, the administrator must use the security recovery utility tool. You cannot unlock an administrative password using the DS Command-Line Interface. The administrative user is forced to establish a new password. Using the **chuser** command, you can specify a password that expires after the initial use, then create a new password. See DS CLI documentation for more information.

Note:

- This task only explains how to use the security recovery utility tool to unlock the Admin account. The topic ″Unlocking a user account″ describes how to unlock a nonadministrative user account.
- The security recovery utility tool is a script that is installed in a file directory. You run the script from the directory.

1. Access the 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 Admin account is unlocked.

Analyzing generic alert traps

The storage unit generates generic alert traps. The traps are sent by an SNMP agent to report asynchronous events.

Generic trap alert identifiers

coldStart (Generic 0)
Issued whenever the SNMP agent is reinitializing. Configuration data has changed.

warmStart (Generic 1)
Issued whenever the SNMP agent is reinitializing. Configuration data might change.

authenticationFailure (Generic 4)
Issued whenever an SNMP message was received but could not be authenticated.

Analyzing service information messages (SIMs) for S/390 and zSeries systems

Service information messages (SIMs) are generated by a storage unit for S/390 and zSeries hosts. Before installation, you can use the customization work sheets to record the following service information: the SIM severity level and how often the stroage unit sends the SIMs (0-9 times) to the console. During installation, either you or the IBM service sales representative must enter this information into the system.

SIM message types

The following SIM message types are generated by the storage unit.

Direct access storage device (DASD) SIM
Tracks DDM failures and problems.

Media SIM
Tracks data check problems on the media.

Storage unit SIM
Tracks storage unit, power control, and other hardware problems.

SIM severity levels

1 acute
An irrecoverable error with possible loss of data. This severity level only applies to DASD SIMs.

2 serious
An irrecoverable error or a data check with loss of access to data.

3 moderate
A system path is not operational and performance might be degraded. This severity level does not apply to media SIMs.

4 service
A recoverable error, equipment checks, or data checks. You can defer repair.

Understanding problem status designations

Each storage unit problem creates a serviceable event. The serviceable event status is either Open or Closed.

Problem status designations

The designations explain where the problem resides in the resolution process.

Open A problem has occurred that requires service. The status will remain in open prior to and during the repair.

Closed
The repair was completed, and that changed the status to closed.

Managing product-specific alert traps

Product-specific alert traps provide information about problems that the storage unit detects and that require corrective action. Either you or the service provider must perform some action for each of these problems. The storage unit generates the following generic, trap-6, product–specific alert traps: 1, 2, 3, 100, 101, 102, 200, 202, 210, 211, 213, 214, 215, 216, 217.

Specific information provided for trap-6, alert trap 1

Each generic, trap-6, product specific alert trap 1 provides the following information.

Problem ID
The identifier assigned by the storage unit for the particular problem that caused the trap.

Presenting time
The time and date of the most recent problem report.

Description
The description of the error condition for the problem ID.

User action
The recommended actions that you need to take to resolve this error condition.

Product-specific alert traps 1 and 2

Product-specific alert traps contain the identifier that the storage unit assigns for the particular problem that caused the trap. The character string *Problemid=n* (where *n* is the problem identifier) is in the description field of the trap. The function for generic, trap-6, product-specific trap 1, and trap 2 on the storage unit is the same. Product-specific alert trap 1 is set as the default.

Product-specific alert traps 1 and 2 generate the following information:

yyyy/mm/dd hh:mm:ss zzz
Provides the time and date of the most recent problem occurrence. (*zzz* represents the time zone.)

data Provides the following detailed information about the trap:

- Storage unit machine type
- Model
- Serial number
- Problem ID
- Severity

partnumber
Provides the 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 ashort 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 a error that requires action, it illuminates one of the message lights on the expansion enclosure. Also informational messages are issued through the storage unit when special events occur.

If your warranty covers the storage unit or you have a service maintenance agreement, you can respond in the following way:

- Contact your authorized service representative.
- Describe the error message, error code, or problem that you have observed.

Note: Your warranty agreement or service maintenance agreement might be affected if you perform your own maintenance.

Use the following topical descriptions of the messages generated by or through the storage unit to help you describe the problem to your service representative.

Managing informational messages from the storage unit

Informational messages are issued through the storage unit as special events occur. Your system administrator determines how these messages should be handled.

Type of informational message

An informational message is issued when your service provider runs the customer-notification diagnostic test. This test verifies that e-mail messages are

being received by those who should receive them. You need to attach your LAN to the DS Network to receive e-mail messages from the storage unit.

Managing the storage unit error messages

The storage unit generates error messages when it detects a situation that requires customer action. This section describes the type of information provided so that you can take the necessary steps to resolve the error condition. In most cases you will need to call your IBM service representative. By providing the listed information your IBM service representative will have an idea where to begin to resolve the problem.

Purpose

The error messages from the storage unit typically contain the following fields.

Product manufacturer ID and date
The ID of the storage unit and the date that it was manufactured.

Storage unit location
The installer enters the storage unit location during the initial installation of the product.

Product machine type and model number
Assigned by IBM at time of manufacturing.

Product serial number
Assigned by IBM at time of manufacturing.

Customer voice phone number
The phone number for customer voice contact.

LMC level of local storage server
The level of the licensed machine code (LMC) of your primary storage unit.

LMC level of remote storage server
The level of the LMC of your secondary or backup storage unit.

Report time and date stamp
The time that this report was generated.

Problem ID
The problem ID that is assigned to this problem by the storage unit. The service provider uses this problem ID to access detailed problem information.

SRN/SRC
A detailed error code that the service provider uses.

Problem status
The problem status state.

Description
A description of the problem.

Additional message
Any additional information that is available.

Reporting resource
The coded resource name that the service provider uses during the repair process.

Failure occurred
The date and time when the failure first occurred.

Last occurrence
The date and time the last occurrence was noted.

Failure count
The number of times that this failure occurred.

Presentation interval
The time between successive e-mail copies of this problem.

Remaining presentations
The number of additional times this e-mail notification will be sent.

Isolation procedure
A pointer to a special procedure in the online service information center.

Failure actions
Actions that the service provider can take.

Probable cause
Information for the service provider.

Failure cause
Information for the service provider.

The following fields are the most useful to you in identifying DDM failures:

- Description
- Reporting resource
- Last occurrence

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 storage enclosure.

Problem

Follow these procedures if you have removed a DDM or storage enclosure that was in use but have received an error message or log entry stating that the DDMs or storage 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 storage 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 storage 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 storage 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 storage 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 a terminal emulator to connect to the server enclosure through the serial port located on the processor card. Use the following terminal emulator settings when you connect through the serial port:

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

2. Use the default user ID of ″guest″ to access the processor card. The `ncnetconf` 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.
3. After the program has started, choose `Configure network parameters` from the ncnetconf Main Menu options.
4. 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 ncnetconf 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 ncnetconf Main Menu.
   c. Select `Apply changes and exit` from the options in the main menu to save your changes and exit the application.
5. To set the IP address, perform the following steps:
   a. Choose `Use static IP address` from the Network configuration menu options.
   b. To change the IP address for this node, choose `IP address for this node` from the Static IP addresses configuration menu options. To change the IP address for the other node, choose `IP address for other node` from the Static IP addresses configuration menu options.

      Note: The ncnetconf program prevents you from setting a network mask IP address that might conflict with the above addresses.
   c. When the `IP Address?` prompt appears, type in the appropriate IP address and press Enter.
   d. Select `Back to Network Configuration` to return to the Network configuration menu.
   e. Select `Back to Main Menu` to return to the ncnetconf Main Menu.
   f. Select `Apply changes and exit` from the options in the main menu to save your changes and exit the application.

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 storage enclosures.

Problem

Follow these procedures if you unsuccessfully attempt to power on or off the server or storage 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 a storage enclosure, push the power button on the rear operator panel of the storage 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 theses 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 storage 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 storage 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 storage 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 storage enclosure has been lost. Replace the specified cable, small-form factor pluggable (SFP), or both, to regain connectivity to the storage 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 storage 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 storage 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 storage 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 storage 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 a storage 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 storage 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

Follow these steps 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 storage 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 storage enclosure that has the system alert indicator lit and repeat this step and the successive steps for that storage 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

Follow these steps 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.

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 storage enclosure.
5. Reconnect the replacement cable to the server or storage 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.

Setting up call home

Use this process to configure notifications.

You must define the customer contact information before configuring notifications.

This task enables you to define Call Home (SMTP), SNMP, and SIMs (service information messages) notifications for a storage unit. See Notification methods for detailed information about these functions.

1. Under **Manage hardware**, select **Storage units**. In Storage units - Main Page, select **Configure notifications** in the **Select Action** drop-down field. Then click **Go**. The Configure notifications - Define Call Home page is displayed.
2. Ensure that **Enable Call Home** is selected to activate Call Home. (This is checked by default.)
3. Complete the SMTP information.
   a. Enter the SMTP server host name.
   b. Enter the SMTP IP address.
   c. Enter the SMTP server ports.
   d. Click **Apply**.
4. Click **Test Call Home connection** to send a connection test and generate a problem log entry. A confirmation message is displayed.
5. Click the **SNMP** tab. The Define SNMP connection page is displayed.
6. Select Enable SNMP notification to define the SNMP connection properties for the selected storage units.
7. Specify either an IP address, a Host name, or both under SNMP trap destination.
8. Specify an SNMP community name of up to 32 characters. This field is used to authenticate requests. 'Public' is selected by default.
9. Optionally specify an SNMP system contact name of up to 32 characters. Enter a destination port. Click **Apply**.
10. Click the **zSeries** tab. The Define Service Information Messages for zSeries page is displayed.
11. Optionally select a SIM severity level in the Severity reporting level for DASD Service Information Messages field.
12. Optionally select a Media Service Information severity level in the Severity reporting level for Media Service Information Messages field.
13. Optionally select a Service Information severity level in the Severity reporting level for a Service Information Messages field.
14. Click Apply.

Sending problem determination data

Use this process to create a problem file and send it to IBM.

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

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 Copy Problem Determination Data page is displayed.
4. Select to either Copy new data or Use existing data. Select if you would like to include Traces, Dumps, or both file types. Click **Next**. The system begins creating a PE package file and displays the Download problem determination data page.

Note: If you want to download data, select the **Download data to local directory** checkbox. The files display as a list of links. To save the files, right click the file link, and select **Save as** to save to a directory on your local drive.

5. Choose to download data to a local directory if you want a copy of the data on the management console. Select the files that contain the problem determination data that you want to download. If you do not want a copy of the problem determination files, do not select the Download data to local directory checkbox. Click **Next** to display the Send problem determination data to IBM page.
6. You must choose to either send all data to IBM or select one or more data files from the table. Click **Next** to display the Verification page.
7. Verify that the information in the table is accurate and click **Finish** to send the selected information to IBM. Use the **Back** button to move through previous pages and modify any incorrect selections in the table.

Contacting IBM

Follow these steps 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, under Real-time Manager, select, in order, Monitor System and 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 problem determination information using ftp

The problem determination information that you collect when you process the mkpe and offloadss commands is forwarded automatically to IBM support personnel using ftp if your storage management console (SMC) is attached to the Internet.

To use the ftp parameter of the mkpe and offloadss commands, you must ensure the following:

- The Internet Raid controllers (IRCs) must be powered on and functioning.
- There must be an active IP connection between your storage management console and both IRCs.

- Adequate disk space must exist on your storage management console (SMC).
- You must be logged into the DS CLI application in interactive mode.

Note: The ftp function works only from your storage management console.

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 statesave data, they also request that you provide the PE package information, as well. There are occasions, however, when only the PE package information is needed for analysis.

Use the following process to provide the requested information to IBM support:

1. Ensure that your SMC meets the prerequisites before issuing the mkpe and offloadss commands.
2. Determine the directory where you want to collect the PE package information. For example, /etc/prepackage
3. Issue the mkpe command from the dscli prompt as follows: dscli>`mkpe -destdir /etc/prepackage IBM.1750-75FA120`. The PE package information is collected in the destination directory and using ftp, is automatically sent to the applicable IBM support URL. When the information has been successfully received by the IBM support site (for example, testcase.boulder.ibm.com), the information files are automatically erased from the collection directory.

   A successful process of this command displays a message like the following:

   ```
   Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
   DS: IBM.1750-75FA120

   The PE package was successfully offloaded to the management node
   ```
4. Determine the directory where you want to collect the statesave information. For example, /etc/statesave
5. Issue the offloadss command as follows (if you have been asked to supply the statesave information): dscli>`offloadss -destdir /etc/statesave IBM.1750-75FA120`. The statesave information is collected in the destination directory and automatically sent to the applicable IBM support URL. When the information has been successfully received by the IBM support site, the information files are automatically erased from the collection directory.

   A successful process of this command displays a message like the following:

   ```
   Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
   DS: IBM.1750-75FA120

   The State Save was successfully offloaded to the management node
   ```

Providing problem determination information when 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 still make use of the mkpe and offloadss commands to collect the information, but you 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 state save data, which

contains additional detailed information. Generally, when IBM support personnel request the state save data, they also request that you provide the PE package information, as well. There are occasions however, when only the PE package information is needed for analysis.

Use the following process to provide the requested information to IBM support:

1. Obtain the address information from IBM Support where you will send your e-mail or CD.
2. Determine the directory where you want to collect the PE package information. For example, /etc/prepackage
3. Issue the mkpe command from the dscli prompt as follows: dscli>`mkpe -destdir /etc/prepackage -noftp IBM.1750-75FA120`. The PE package information is collected in the destination directory.
4. Create a zip file of the information if there are many files or copy the files to a CD.
5. 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.
6. Determine the directory where you want to collect the state save information. For example, /etc/statesave
7. Issue the offloadss command from the dscli prompt as follows: dscli>`offloadss -destdir /etc/statesave -noftp IBM.1750-75FA120`. The state save information is collected in the destination directory.
8. Create a zip file of the information or copy the files to a CD.
9. Create and send an e-mail with your attached state save file, to the address that IBM Support has directed your file be sent. Or, send your CD by mail to the address that IBM Support has given you.
10. Wait for confirmation of receipt from IBM Support and copy your state save 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.
11. Delete the PE package files from your collection directory.
12. Delete the state save files from your collection directory.

Chapter 15. 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 in preparation for a possible disaster by backing up, copying, and mirroring your data at the production and recovery 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. This means that you must be able to recover your data at a backup (normally your recovery site).

In an event of an planned event or unplanned disaster, you can use failover and failback modes as part of your recovery solution. Failover and failback modes help in reducing the time that is required to synchronize remote mirror and copy volumes after switching between the production (local) and the recovery sites during planned and unplanned outages. A failover is the process of temporarily switching production to a recovery site following a planned outage, such as a scheduled maintenance period or an unplanned outage, such as a disaster. A failback operation is the process of returning production to its original location.

Using failover and failback modes for planned and unplanned outages: Most users decide to use a combination of remote mirror and copy and point-in-time copy (FlashCopy) features to form a comprehensive enterprise solution for disaster recovery. Recovery procedures can include failover and failback modes with the following remote mirror and copy functions and FlashCopy to help reduce the time it takes to recover at the recovery site. The main difference is how data consistency is achieved at your recovery site.

- Metro Mirror, Global Copy
- Global Mirror, which also uses Global Copy

The data consistency differences are outlined below:

Manual or external software required (without Global Mirror)
If you use Metro Mirror, Global Copy, and FlashCopy to create a consistent and restartable copy at your recovery site, you must do a manual and periodic suspend at your local site. This means using *freeze and run* commands or using external automated software to create data consistency. (Automation software is not provided with the storage unit; it must be supplied by the user.)

Automated (with Global Mirror)
If you use Global Mirror, the process to create a consistent and restartable copy at your recovery site is done using an automated process, with minimal or no interruption to your applications. A master storage unit (along with subordinate storage units) control the creation of consistency groups, which can be created many times per hour to increase the currency of data that is captured in the consistency groups at the recovery 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 when creating consistent copies of data. The formation of these consistency groups is coordinated by the master storage unit, which sends commands over remote mirror and copy links to its subordinate storage units.

If you have a disaster at your local site and have to start production at your recovery site, you can use the consistent point-in-time data from the consistency group at your recovery site to recover during a failback operation.

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 failover to your recovery site to continue operations. After your production site is operational, you can perform a subsequent failback operation to move production back to its original location.

Moving production to Site B after planned outages (failover)

When you schedule a planned outage at your production site (Site A), you can switch production to your recovery site (Site B), allowing the processing of data to resume at Site B. This process is known as a failover recovery.

The storage units at both Site A and Site B must be functional and reachable.

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

For this scenario, assume that all I/O to Site A has ceased because of a planned outage, such as a scheduled maintenance. The failover operation is issued to the storage unit that will become the primary. That is, production will failover to Site B during this outage, which will make the target volumes at Site B convert to source volumes and cause them to suspend. Your original source volumes at Site A remain in the state they were in at the time of the site switch. See Table 3 on page 219 for 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 will change a target volume into a suspended source volume, while leaving the source volume in its current state.

The following assumptions are made for this scenario:

- Applications are updating the source volumes located at Site A.
- Paths are established from Site A to Site B.
- Volume pairs are in duplex state.

Perform the following steps using the DS Storage Manager. You can also use the DS CLI to perform Copy Services functions.

1. When the planned outage window is reached, quiesce applications to cease all write I/O from updating the source volumes. Quiescing your applications should occur as part of a planned outage and the outage should be brief. Depending on the host operating system, it might be necessary to unmount the source volumes.

2. Perform a failover recovery operation to Site B. After the failover operation has processed successfully, the volumes at Site B transition from target to source volumes.
3. Depending on your path design and when the source storage unit becomes available, create paths in the opposite direction from Site B to Site A. You will need the paths in the opposite direction if you want to transfer the updates back to Site A.
4. Depending on your operating system, it might be necessary to rescan your fibre-channel connected devices (to remove device objects for the site A volumes and recognize the new source volumes).
5. Mount your target volumes (now the new source volumes) on the target storage unit at Site B.
6. Start all applications. After the applications start, all write I/O operations to the source volumes are tracked. Depending on your plans regarding Site A, the volume pairs can remain suspended (if you want to do offline maintenance).
7. When your scheduled maintenance is complete, schedule a time for a failback recovery operation to initiate the transfer of data back to Site A. This process resynchronizes the volumes at Site A with the volumes at Site B.

Note: 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 3. 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						
Disaster at Site A	Failover	Site B	Volume B, Volume A	Volume B1 : Volume A1	Volume A1 -> Volume B1 (Suspended) The volume pair might display as full or pending duplex state if host write operations have stopped.	Volume B1 -> Volume A1 (Suspended)
2 (Site A volumes must be in suspended state)						
Return production to Site A	Failback	Site A	Volume A, Volume B	Volume A1 : Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1
3a (Site B volumes should be in a suspended state)						

Table 3. 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
Return to production (Site B) **Note:** If Site A is still not operational; production can continue at Site B.	Failback	Site B	Volume B, Volume A	Volume B1 : Volume A1	Volume B1 -> Volume A1	Volume B1 -> Volume A1
3b (prepare to return to production (Site A) from production (Site B)	Failover	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume B1 -> Volume A1 (Suspended state; the volume pair might display full or pending state if host write operations have stopped.)
3c (Site A volumes should be in a suspended state)						
Return to production - Site A	Failback	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1

Moving production to Site B after unplanned outages (failover)

In this situation, you have lost access to your production site, Site A. You can switch production to your recovery site (Site B) allowing the processing of data to resume at Site B. This process is known as a failover recovery.

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

For this scenario, assume that an unexpected failure occurs at Site A, making it unavailable. Production must be moved to Site B.

The failover operation is issued to the storage unit that will become the primary. That is, production will failover to Site B during this outage, which will make the target volumes at Site B convert to source volumes and cause them to suspend. Your original source volumes at Site A remain in the state they were in at the time of the site switch. When Site A is available again, application I/O is switched back from Site B to Site A. See Table 4 on page 221 for an example of the implementation of failover and failback operations.

Perform the following steps using the DS Storage Manager. You can also use the DS CLI to perform Copy Services functions.

1. Perform a failover recovery operation to Site B. After the failover operation has processed successfully, the volumes at Site B transition from target to source volumes.
2. Mount your target volumes on your server at Site B.
3. Start your applications on your server at Site B.
4. After Site A recovers, proceed with the following steps, which are the first taken to recover the volumes at Site A.
   a. Create paths between LSSs at Site B to Site A to allow the volumes at Site A to be synchronized with the Site B volumes.
   b. Delete any remote mirror and copy volume relationships that still exist from the source volumes.
   c. Wait until the volumes are in full duplex state, then schedule a time to perform a failback recovery operation using the volumes at Site A. This process resynchronizes the volumes at Site A with the volumes at Site B.

      Note: Failback recovery operations are usually used after a failover recovery has been issued to restart mirroring either in the reverse direction (remote site to local site) or original direction (local site to remote site).

Table 4. Failover and failback implementation

Step	Operation	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						
Disaster at Site A	Failover	Site B	Volume B, Volume A	Volume B1 : Volume A1	Volume A1 -> Volume B1 (Suspended) The volume pair might display as full or pending duplex state if host write operations have stopped.	Volume B1 -> Volume A1 (Suspended)
2 (Site A volumes must be in suspended state)						
Return production to Site A	Failback	Site A	Volume A, Volume B	Volume A1 : Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1
3a (Site B volumes should be in a suspended state)						
Return to production (Site B) **Note:** If Site A is still not operational; production can continue at Site B.	Failback	Site B	Volume B, Volume A	Volume B1 : Volume A1	Volume B1 -> Volume A1	Volume B1 -> Volume A1

Table 4. Failover and failback implementation (continued)

Step	Operation	MC connectivity required to	Format of source volume and target volume	Format of source and target volume pair	Result: Site A	Result: Site B
3b (prepare to return to production (Site A) from production (Site B)	Failover	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume B1 -> Volume A1 (Suspended state; the volume pair might display full or pending state if host write operations have stopped.)
3c (Site A volumes should be in a suspended state)						
Return to production - Site A	Failback	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1

Returning production to 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 5 on page 223 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. 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.

Table 5. 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						
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 a full or pending duplex state if host write operations have stopped.	Volume B1 -> Volume A1 (Suspended)

Table 5. 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
2 (Site A volumes must be in suspended state)						
Return production to Site A	Failback	Site A	Volume A, Volume B	Volume A1 : Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1
3a (Site B volumes should be in a suspended state)						
Return to production (Site B) **Note:** If Site A is still not operational; production can continue at Site B.	Failback	Site B	Volume B, Volume A	Volume B1 : Volume A1	Volume B1 -> Volume A1	Volume B1 -> Volume A1
3b (prepare to return to production (Site A) from production (Site B)	Failover	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume B1 -> Volume A1 (Suspended state; the volume pair might display a full or pending state if host write operations have stopped.)
3c (Site A volumes should be in a suspended state)						
Return to production - Site A	Failback	Site A	Volume A, Volume B	Volume A1: Volume B1	Volume A1 -> Volume B1	Volume A1 -> Volume B1

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

Use this process to set up your environment to use Global Mirror for disaster recovery. Global Mirror asynchronously copies data from a host to a recovery site and maintains data on a storage unit at the recovery site.

You can configure the following Global Mirror environment, which uses two sites (local and remote) and three volumes (volume A, volume B, and volume C) on two or more storage units.

Local site
Contains volume A (the source volume), which is copied to the recovery site using Global Copy

Recovery (or remote) site
Contains volume B (the target volume and FlashCopy source volume) and volume C (the FlashCopy target volume)

Note: The storage unit at the local site is designated as the Global Mirror master and all other local (or production) storage units are designated as subordinate storage units. The master 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.

Perform the following steps to set up your environment to use Global Mirror and 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 individuals tasks, see the related topics section.

1. Quiesce host I/O that is updating the primary volumes that is located in the local site.
2. Check the status of Global Mirror to ensure that Global Mirror is in a good state and that a consistency group has formed successfully at the desired interval.
3. Wait until two consistency groups form successfully.
4. Pause Global Mirror and suspend the Global Copy pairs (A->B).
5. Create paths from the B volumes to the A volumes.
6. Perform recovery failover on the B volumes (B->A).
7. Perform the fast reverse restore process from the C volumes to the B volumes, selecting the **Initiate background copy** option.

Note:

a. When you initiate the fast reverse restore process, Volume C becomes unusable.

b. There must be *no* I/O allowed to the B or C volumes during the fast reverse restore process.

8. Wait for the background copy to complete before continuing to the next step. The C to B FlashCopy relationship ends when the background copy completes.
9. Initiate the FlashCopy from Volume B back to Volume C. Ensure that you also select the **Enable Change Recording** and **Inhibit writes to target volume** options. This creates a backup copy of the consistency group before applications begin to update the B volumes.
10. Start the host I/O at the remote site on the B volumes. You will remain in production on the remote site in this configuration until you are ready to return production to the local site.
11. When you are ready to return production to the local site, perform recovery failback (B->A) to resynchronize the A volumes. The application at the remote site remains active.
12. After the resynchronize process has completed its first pass, quiesce the applications at your remote site so that the resynchronization can complete.
13. When the resynchronization has completed (no out-of-sync tracks), perform recovery failover and failback with Global Copy on Volume A to re-create the Volume A to Volume B Metro Mirror relationship.
14. Start the host I/O at the local site on the A volumes.
15. Resume the Global Mirror process.

Using Global Mirror for an unplanned failover and failback

Use this process to manage an unplanned failover and failback that involves two sites which are referred to as the local (or production) and remote (or recovery) sites.

Global Mirror provides two-site extended distance remote copy disaster recovery. When a disaster occurs at the local site, you must initiate the failover and failback recovery of consistent data on the remote site. Host activity can resume on the local site when the host recovers but not before a consistent set of data is copied to all primary volumes on the local site.

With Global Mirror, the data that the host writes to the storage unit at the local site is asynchronously shadowed to the storage unit at the remote site. A consistent copy of the data is then automatically maintained on the storage unit at the remote site.

Use of Global Mirror does not guarantee against data loss. During a disaster, data can only be restored to the last known consistent increment that was created. This means that data that is written to the primary site and is waiting to be transferred to the secondary site is lost whenever the two storage units can no longer communicate.

The following considerations apply when you use the Global Mirror recovery process:

- The Global Mirror master might still be running at the local site, especially if the disaster at the local site is a rolling disaster in which not all components fail simultaneously.

- The consistent copy at the remote site is *not* the secondary volume, but it is the FlashCopy target whose source is the secondary volume.
- Formation of a consistency group might have been in progress at the time of the failure.
- You can speed up recovery processing by choosing the **Fast Reverse** restore process that is explained later in this section.

Perform the following steps to use Global Mirror for an unplanned failover using the DS Storage Manager. You can also use the DS CLI for these steps.

Note: For details on individuals 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.

Accessibility

Accessibility features provide users who have disabilities with the ability to successfully access information and use technology.

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features

These are the major accessibility features in the IBM TotalStorage 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 TotalStorage 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 TotalStorage 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 for downloading and printing publications

Permissions for the use of the information you have selected for download are granted subject to the following terms and conditions and your indication of acceptance thereof.

Personal Use: You may reproduce this information for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of this information, or any portion thereof, without the express consent of the manufacturer.

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

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

The manufacturer reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the information is detrimental to its interest or, as determined by the manufacturer, the above instructions are not properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations. THE MANUFACTURER MAKES NO GUARANTEE ABOUT THE CONTENT OF THIS INFORMATION. THE INFORMATION IS PROVIDED ″AS-IS″ AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

By downloading or printing information from this site, you have indicated your agreement with these terms and conditions.

Trademarks

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

- AIX
- DB2
- DFSMS/MVS
- DFSMS/VM
- e (logo)
- Enterprise Storage Server
- ES/9000
- ESCON
- FICON
- FlashCopy
- HACMP
- IBM
- Intellistation
- MVS/ESA
- Netfinity
- NetVista
- Operating System/400
- OS/400
- RS/6000
- S/390
- Seascape
- SNAP/SHOT
- SP
- System/390
- TotalStorage
- Versatile Storage Server
- Virtualization Engine
- VSE/ESA
- z/Architecture
- z/OS
- z/VM
- zSeries

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

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

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

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

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

Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.

Federal Communications Commission (FCC) statement

This equipment has been tested and complies with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the users authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Industry Canada compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conform à la norme NMB-003 du Canada.

European community compliance statement

This product is in conformity with the protection requirements of EC Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

Germany only

Zulassungsbescheinigung laut Gesetz ueber die elektromagnetische 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 TotalStorage and other Resiliency Family products.

This glossary includes selected terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172–1990, copyright 1990 by the American National Standards Institute (ANSI), 11 West 42nd Street, New York, New York 10036. Definitions derived from this book have the symbol (A) after the definition.
- *IBM Terminology*, which is available online at the following Web site: http://www-306.ibm.com/ibm/terminology/index.html. Definitions derived from this source have the symbol (GC) after the definition.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions derived from this book have the symbol (I) after the definition. Definitions taken from draft international standards, committee drafts, and working papers that the ISO/IEC JTC1/SC1 is developing have the symbol (T) after the definition, indicating that final agreement has not been reached among the participating National Bodies of SC1.

This glossary uses the following cross-reference forms:

See Refers the reader to one of two kinds of related information:

- A term that is the expanded form of an abbreviation or acronym. This expanded form of the term contains the full definition.
- A synonym or more preferred term

See also
Refers the reader to one or more related terms.

Contrast with
Refers the reader to a term that has an opposite or substantively different meaning.

Numerics

750 A model of the Enterprise Storage Server featuring a 2-way processor with limited physical storage capacity. This model can be updated to the model 800.

800 A model of the Enterprise Storage Server featuring a standard processor or an optional Turbo processor. The Model 800 supports RAID 5, RAID 10, and 15000 rpm drives. Model 800 supersedes Model F20.

1750 The machine type for the IBM TotalStorage 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 TotalStorage DS8000 series. Models for the DS8000 series include the base units 921, 922 and 9A2 and expansion units 92E and 9AE.

3390 The machine number of an IBM disk storage system. The Enterprise Storage Server, when interfaced to IBM zSeries hosts, is set up to appear as one or more 3390 devices, with a choice of 3390-2, 3390-3, or 3390-9 track formats.

3990 The machine number of an IBM control unit.

7133 The machine number of an IBM disk storage system. The Model D40 and 020 drawers of the 7133 can be installed in the 2105-100 expansion enclosure of the ESS.

A

access
1) To obtain the use of a computer resource.

2) In computer security, a specific type of interaction between a subject and an object that results in flow of information from one to the other.

access-any mode
One of the two access modes that can be set for the storage unit during initial configuration. It enables all fibre-channel-attached host systems with no defined access profile to access all logical volumes on the storage unit. With a profile defined in DS Storage Manager for a particular host, that host has access only to volumes that are assigned to the WWPN for that host. See also *pseudo host* and *worldwide port name*.

ACK See *request for acknowledgment and acknowledgment*.

agent A program that automatically performs some service without user intervention or on a regular schedule. See also *subagent*.

alert A message or log that a storage unit generates as the result of error event collection and analysis. An alert indicates that a service action is required.

allegiance
For zSeries, a relationship that is created between a device and one or more channel paths during the processing of certain conditions. See also *implicit allegiance*, *contingent allegiance*, and *reserved allegiance*.

allocated storage
The space that is allocated to volumes but not yet assigned. Contrast with *assigned storage*.

American National Standards Institute (ANSI)
An organization of producers, consumers, and general interest groups that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States. (A)

Anonymous
In the DS Storage Manager, the label on an icon that represents all connections that are using fibre-channel adapters between the storage unit and hosts but are not completely defined to the storage unit. See also *anonymous host*, *pseudo host*, and *access-any mode*.

anonymous host
Synonym for *pseudo host*. Contrast with *Anonymous* and *pseudo host*.

ANSI See *American National Standards Institute*.

APAR See *authorized program analysis report*. (GC)

API See *application programming interface*.

application programming interface
An interface that allows an application program that is written in a high-level language to use specific data or functions of the operating system or another program.

arbitrated loop
A fibre-channel topology that enables the interconnection of a set of nodes. See also *point-to-point connection* and *switched fabric*.

array An ordered collection, or group, of physical devices (disk drive modules) that is used to define logical volumes or devices. In the storage unit, an array is a group of disks that the user designates to be managed by the RAID technique. See also *redundant array of independent disks*.

ASCII (American National Standard Code for Information Interchange) The standard code, using a coded character set consisting of 7-bit coded characters (8 bits including parity check), that is used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. (A) Some organizations, including IBM, have used the parity bit to expand the basic code set.

assigned storage
The space that is allocated to a volume and that is assigned to a port.

authorized program analysis report (APAR)
A request for correction of a defect in a current release of an IBM-supplied program. (GC)

availability
The degree to which a system or resource is capable of performing its normal function. See *data availability*.

B

bay The physical space that is used for installing SCSI, ESCON, and fibre-channel host adapter cards. The DS6000 has four bays, two in each cluster. See also *service boundary*.

bit 1) Either of the digits 0 or 1 when used in the binary numeration system. (T)

2) The storage medium required to store a single binary digit. See also *byte*.

block A string of data elements recorded or transmitted as a unit. The elements may be characters, words, or physical records. (T)

A group of consecutive bytes used as the basic storage unit in fixed-block architecture (FBA). All blocks on the storage device are the same size (fixed size). See also *fixed-block architecture* and *data record*.

byte 1) A group of eight adjacent binary digits that represent one EBCDIC character.

2) The storage medium required to store eight bits. See also *bit*.

C

cache A special-purpose buffer storage, smaller and faster than main storage, used to hold a copy of instructions and data obtained from main storage and likely to be needed next by the processor. (T)

cache fast write
A form of the fast-write operation in which the storage server writes the data directly to cache, where it is available for later destaging.

cache hit
An event that occurs when a read operation is sent to the cluster, and the requested data is found in cache. Contrast with *cache miss*.

cache memory
Memory, typically volatile memory, that a storage server uses to improve access times to instructions or data. The cache memory is typically smaller and faster than the primary memory or storage medium. In addition to residing in cache memory, the same data also resides on the storage devices in the storage unit.

cache miss
An event that occurs when a read operation is sent to the cluster, but the data is not found in cache. Contrast with *cache hit*.

call home
A communication link established between the DS6000 and a service provider. The DS6000 can use this link to place a call to IBM or to another service provider when it requires service. With access to the machine, service personnel can perform service tasks, such as viewing error logs and problem logs or initiating trace and dump retrievals. See also *heartbeat* and *remote technical assistance information network*.

cascading
1) Connecting network controllers to each other in a succession of levels to concentrate many more lines than a single level permits.

2) In high-availability cluster multiprocessing (HACMP), cascading pertains to a cluster configuration in which the cluster node with the highest priority for a particular resource acquires the resource if the primary node fails. The cluster node relinquishes the resource to the primary node upon reintegration of the primary node into the cluster.

catcher
A server that service personnel use to collect and retain status data that an DS6000 sends to it.

CCR See *channel command retry*.

CCW See *channel command word*.

CD See *compact disc*.

CEC See *computer-electronic complex*.

channel
For zSeries, the part of a channel subsystem that manages a single I/O interface between a channel subsystem and a set of control units.

channel command retry (CCR)
For zSeries, the protocol used between a channel and a control unit that enables the control unit to request that the channel reissue the current command.

channel command word (CCW)
For zSeries, a data structure that specifies an I/O operation to the channel subsystem.

channel path
For zSeries, the interconnection between a channel and its associated control units.

channel subsystem
For zSeries, the part of a host computer that manages I/O communication between the program and any attached control units.

channel-subsystem image
For zSeries, the logical functions that a system requires to perform the function of a channel subsystem. With ESCON multiple image facility (EMIF), one channel subsystem image exists in the channel subsystem for each logical partition (LPAR). Each image appears to be an independent channel subsystem program, but all images share a common set of hardware facilities.

CKD See *count key data*.

CLI See *command-line interface*. See also *Copy Services command-line interface*.

cluster
1) A partition capable of performing all DS6000 functions. With two clusters in the DS6000, any operational cluster can take over the processing of a failing cluster.

2) In the AIX operating system, a group of nodes within a complex.

cluster processor complex (CPC)
The unit within a cluster that provides the management function for the DS6000. It consists of cluster processors, cluster memory, and related logic.

command-line interface (CLI)
An interface provided by an operating system that defines a set of commands and enables a user (or a script-like language) to issue these commands by typing text in response to the command prompt (for example, DOS commands or UNIX shell commands). See also *Copy Services command-line interface*.

compact disc
An optically read disc, typically storing approximately 660 MB. CD-ROM (compact disc read-only memory) refers to the read-only format used to distribute 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)

computer-electronic complex (CEC)
The set of hardware facilities associated with a host computer.

concurrent copy
A facility on a storage server that enables a program to make a backup of a data set while the logical volume remains available for subsequent processing. The data in the backup copy is frozen at the point in time that the server responds to the request.

concurrent installation of licensed internal code Process of installing licensed internal code on a 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 configuration of the input/output (I/O) subsystem through the user interface that the storage unit provides for this function.

consistency group
A group of volumes participating in FlashCopy relationships in a logical subsystem, across logical subsystems, or across multiple storage units that must be kept in a consistent state to ensure data integrity.

consistency group interval time
The value in seconds that indicates the length of time between the formation of consistency groups.

consistent copy
A copy of a data entity (a logical volume, for example) that contains the contents of the entire data entity at a single instant in time.

console
A user interface to a server, for example, the interface provided on a personal computer. See also *IBM TotalStorage Management Console*.

contingent allegiance
For zSeries, a relationship that is created in a control unit between a device and a channel when the channel accepts unit-check status. The allegiance causes the control unit to guarantee access; the control unit does not present the busy status to the device. The allegiance enables the channel to retrieve sense data that is associated with the unit-check status on the channel path associated with the allegiance.

control path
The route that is established from the master storage unit to the subordinate storage unit when more than one storage unit participates in a Global Mirror session. If there is only one storage unit (the master) in the Global Mirror session, no control path is required.

control unit (CU)
1) A device that coordinates and controls the operation of one or more input/output devices, and synchronizes the operation of such devices with the operation of the system as a whole.

2) For zSeries, a storage server with ESCON or OEMI interfaces. The control unit adapts a native device interface to an I/O interface that a zSeries host system supports.

3) The portion of the storage unit that supports the attachment of emulated count key data devices over ESCON, FICON, or OEMI interfaces. See also *cluster*.

control-unit image
For zSeries, a logical subsystem that is accessed through an ESCON I/O interface. One or more control-unit images exist in each control unit. Each image appears as an independent control unit, but all control-unit images share a common set of hardware facilities. The DS6000 can emulate 3990-3, TPF, 3990-6, or 2105 control units.

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 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.

count field
The first field of a count key data (CKD) record. This eight-byte field contains a four-byte track address (CCHH). It defines the cylinder and head that are associated with the track, and a one-byte record number (R) that identifies the record on the track. It defines a one-byte key length that specifies the length of the record's key field (0 means no key field). It defines a two-byte data length that specifies the length of the record's data field (0 means no data field). Only the end-of-file record has a data length of zero.

count key data (CKD)
For zSeries, a data-record format employing self-defining record formats in which each record is represented by up to three fields: a *count* field that identifies the record and specifies its format, an optional *key* field that identifies the data area contents, and an optional *data* field that typically contains the user data. For CKD records on the storage unit, the logical volume size is defined in terms of the device emulation mode (3390 or 3380 track format). The count field is always 8 bytes long and contains the lengths of the key and data fields, the key field has a length of 0 to 255 bytes, and the data field has a length of 0 to 65 535 or the maximum that will fit on the track. See also *data record*.

CPC See *cluster processor complex*.

CRC See *cyclic redundancy check*.

CU See *control unit*.

CUIR See *control-unit initiated reconfiguration*.

custom volume
A volume in count-key-data (CKD) format that is not a standard volume, which means that it does not necessarily present the same number of cylinders and capacity to its assigned logical control unit as provided by one of the following standard zSeries volume types: 3390-2, 3390-3, 3390-9, 3390-2 (3380-track mode), or 3390-3 (3380-track mode). See also *count-key-data*, *interleave*, *standard volume*, and *volume*.

CUT See *Coordinated Universal Time*.

cyclic redundancy check (CRC)
A redundancy check in which the check key is generated by a cyclic algorithm. (T)

cylinder
A unit of storage on a CKD device with a fixed number of tracks.

D

DA See *device adapter*. See also *IBM Serial Storage adapter*.

daisy chain
See *serial connection*.

DASD See *direct access storage device*.

DASD fast write (DFW)
A function of a storage server in which active write data is stored in nonvolatile cache, thus avoiding exposure to data loss.

data availability
The degree to which data is available when needed, typically measured as a percentage of time that the system would be capable of responding to any data request (for example, 99.999% available).

data compression
A technique or algorithm used to encode data such that the encoded result can be stored in less space than the original data. The original data can be recovered from the encoded result through a reverse technique or reverse algorithm. See also *compression*.

Data Facility Storage Management Subsystem (DFSMS)
An operating environment that helps automate and centralize the management of storage. To manage storage, DFSMS provides the storage administrator with control over data class, storage class, management class, storage group, and automatic class selection routine definitions.

data field
The optional third field of a count key data (CKD) record. The count field specifies the length of the data field. The data field contains data that the program writes.

data record
The basic unit of zSeries storage on a 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 address
For zSeries, the field of an ESCON device-level frame that selects a specific device on a control-unit image.

device ID
The unique two-digit hexadecimal number that identifies the logical device.

device interface card
A physical subunit of a storage cluster that provides the communication with the attached device drive modules.

device number
For zSeries, a four-hexadecimal-character identifier, for example 13A0, that the systems administrator associates with a device to facilitate communication between the program and the host operator. The device number is associated with a subchannel.

device sparing
A subsystem function that automatically copies data from a failing device drive module to a spare device drive module. The subsystem maintains data access during the process.

DFS See *distributed file service*.

DFSMS
See *Data Facility Storage Management Subsystem*.

direct access storage device (DASD)
1) A mass storage medium on which a computer stores data.

2) A disk device.

disk cage
A container for disk drives. Each disk cage supports eight disk packs (64 disks).

disk drive
Standard term for a disk-based nonvolatile storage medium. The 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 TotalStorage 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 TotalStorage DS Storage Manager*.

duplex
1) Regarding Copy Services, the state of a volume pair after Remote Mirror and Copy has completed the copy operation and the volume pair is synchronized.

2) In general, pertaining to a communication mode in which data can be sent and received at the same time.

dynamic sparing
The ability of a storage server to move data from a failing disk drive module (DDM) to a spare DDM while maintaining storage functions.

E

E10 The predecessor of the F10 model of the Enterprise Storage Server. See also *F10*.

E20 The predecessor of the F20 model of the Enterprise Storage Server. See also *F20*.

EBCDIC
See *extended binary-coded decimal interchange code*.

EC See *engineering change*.

ECKD See *extended count key data*.

eight pack
See *disk pack*.

electrostatic discharge (ESD)
An undesirable discharge of static electricity that can damage equipment and degrade electrical circuitry.

emergency power off (EPO)
A means of turning off power during an emergency, usually a switch.

EMIF See *ESCON multiple image facility*.

enclosure
A unit that houses the components of a storage subsystem, such as a control unit, disk drives, and power source.

end of file
A coded character recorded on a data medium to indicate the end of the medium. On a count-key-data direct access storage device, the subsystem indicates the end of a file by including a record with a data length of zero.

engineering change (EC)
An update to a machine, part, or program.

Enterprise Systems Architecture/390 (ESA/390)
An IBM architecture for mainframe computers and peripherals. Processor systems that follow the ESA/390 architecture include the ES/9000® family. See also *z/Architecture*.

Enterprise Systems Connection (ESCON)
1) A zSeries computer peripheral interface. The I/O interface uses zSeries logical protocols over a serial interface that configures attached units to a communication fabric.

2) A set of IBM products and services that provide a dynamically connected environment within an enterprise.

EPO See *emergency power off*.

ERDS See *error-recording data set*.

ERP See *error recovery procedure*.

error-recording data set (ERDS)
On zSeries hosts, a data set that records data-storage and data-retrieval errors. A service information message (SIM) provides the error information for the ERDS.

error recovery procedure (ERP)
Procedures designed to help isolate and, where possible, to recover from errors in equipment. The procedures are often used in conjunction with programs that record information on machine malfunctions.

ESA/390
See *Enterprise Systems Architecture/390*.

ESCD See *ESCON director*.

ESCON
See *Enterprise System Connection*.

ESCON channel
A zSeries channel that supports ESCON protocols.

ESCON director (ESCD)
An I/O interface switch that allows the interconnection of multiple ESCON interfaces in a distributed-star topology.

ESCON host systems
zSeries hosts that attach to the DS6000 with an ESCON adapter. Such host systems run on operating systems that include MVS, VSE, TPF, or versions of VM.

ESCON multiple image facility (EMIF)
For zSeries, a function that enables LPARs to share an ESCON channel path by providing each LPAR with its own channel-subsystem image.

EsconNet
In the DS Storage Manager, the label on a pseudo host icon that represents a host connection that uses the ESCON protocol and that is not completely defined on the DS6000. See also *pseudo host* and *access-any mode*.

ESD See *electrostatic discharge*.

eserver
See *IBM e(logo)server*.

ESSNet
See *IBM TotalStorage Enterprise Storage Server Network*.

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 binary-coded decimal interchange code (EBCDIC)
An IBM-developed coding scheme used to represent various alphabetic, numeric, and special symbols with a coded character set of 256 eight-bit codes.

extended count key data (ECKD)
An extension of the count key data (CKD) architecture.

extent A continuous space on a disk that is occupied by or reserved for a particular data set, data space, or file. The unit of increment is a track. See also *multiple allegiance* and *parallel access volumes*.

extent pool
A groups of extents. See also *extent*.

F

fabric In fibre-channel technology, a routing structure, such as a switch, receives addressed information and routes to the appropriate destination. A fabric can consist of more than one switch. When multiple fibre-channel switches are interconnected, they are said to be *cascaded*.

failback
Pertaining to a cluster recovery from failover following repair. See also *failover*.

failover
Pertaining to the process of transferring all control to a single cluster when the other cluster in the storage unit fails. See also *cluster* and *failback*.

fast write
A write operation at cache speed that does not require immediate transfer of data to a disk drive. The subsystem writes the data directly to cache, to nonvolatile storage, or to both. The data is then available for destaging. A fast-write operation reduces the time an application must wait for the I/O operation to complete.

FBA See *fixed-block architecture*.

FC See *feature code*. **Note**: *FC* is a common abbreviation for fibre channel in the industry, but the DS6000 customer documentation library reserves *FC* for feature code.

FC-AL See *Fibre Channel-Arbitrated Loop*.

FCP See *fibre-channel protocol*.

FCS See *fibre-channel standard*.

feature code (FC)
A code that identifies a particular orderable option and that is used by service personnel to process hardware and software orders. Individual optional features are each identified by a unique feature code.

fibre channel
A data-transmission architecture based on the ANSI Fibre Channel standard, which supports full-duplex communication. The DS6000 supports data transmission over fiber-optic cable through its fibre-channel adapters. See also *fibre-channel protocol* and *fibre-channel standard*.

Fibre Channel-Arbitrated Loop (FC-AL)
An implementation of the Fibre Channel standard that uses a ring topology for the communication fabric. Refer to American National Standards Institute (ANSI) X3T11/93-275. In this topology, two or more fibre-channel end points are interconnected through a looped interface.

fibre-channel connection (FICON)
A fibre-channel communications protocol designed for IBM mainframe computers and peripherals.

fibre-channel protocol (FCP)
A protocol used in fibre-channel communications with five layers that define how fibre-channel ports interact through their physical links to communicate with other ports.

Fibre-Channel standard (FCS)
An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric. The protocol has two layers. The IP layer defines basic interconnection protocols. The upper layer supports one or more logical protocols (for example, FCP for SCSI command protocols and SBCON for zSeries command protocols). Refer to

American National Standards Institute (ANSI) X3.230-199x. See also *fibre-channel protocol*.

fibre-channel topology
An interconnection topology supported on fibre-channel adapters. See also *point-to-point connection*, *switched fabric*, and *arbitrated loop*.

FICON
See *fibre-channel connection*.

FiconNet
In the DS Storage Manager, the label on a pseudo host icon that represents a host connection that uses the FICON protocol and that is not completely defined on the 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 10^9 bytes. A gigabyte of memory is 2^{30} 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 DS6000 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 Multi-Processing*.

hard disk drive (HDD)
1) A storage medium within a storage server used to maintain information that the storage server requires.

2) A mass storage medium for computers that is typically available as a fixed disk (such as the disks used in system units of personal computers or in drives that are external to a personal computer) or a removable cartridge.

hardware service manager (HSM)
An option on an AS/400 or iSeries host that enables the user to display and work with system hardware resources and to debug input-output processors (IOP), input-output adapters (IOA), and devices.

HCD See *Hardware Configuration Data*.

HDA See *head and disk assembly*.

HDD See *hard disk drive*.

hdisk An AIX term for storage space.

head and disk assembly (HDA)
The portion of an HDD associated with the medium and the read/write head.

heartbeat
A status report sent at regular intervals from the 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 Multi-Processing (HACMP)
Software that provides host clustering, so that a failure of one host is recovered by moving jobs to other hosts within the cluster.

high-speed link (HSL)
A hardware connectivity architecture that links system processors to system input/output buses and other system units.

home address
A nine-byte field at the beginning of a track that contains information that identifies the physical track and its association with a cylinder.

hop Interswitch connection. A hop count is the number of connections that a particular block of data traverses between source and destination. For example, data traveling from one hub over a wire to another hub traverses one hop.

host See *host system*.

host adapter
A physical subunit of a storage server that provides the ability to attach to one or more host I/O interfaces.

host name
The Internet address of a machine in the network. The host name can be entered in the host definition as the fully qualified domain name of the attached host system, such as `mycomputer.city.company.com`, or as the subname of the fully qualified domain name, for example, `mycomputer`. See also *host system*.

host processor
A processor that controls all or part of a user application network. In a network, the processing unit in which the data communication access method resides. See also *host system*.

host system
A computer, either of the mainframe (for example, zSeries) or of the open-systems type, that is connected to theDS6000. zSeries hosts are connected to the DS6000 through ESCON interfaces. Open-systems hosts are connected to the DS6000 by SCSI or fibre-channel interfaces.

hot plug
Pertaining to the ability to add or remove a hardware facility or resource to a unit while power is on.

HSL See *high-speed link*.

HSM See *hierarchical storage management* or *Hardware Service Manager*.

I

i5/OS The IBM operating system that runs the IBM i5/OS and e(logo)Server i5 server families of servers.

IBM e(logo)server
The IBM brand name for a series of server products that are optimized for e-commerce. The products include the iSeries, pSeries, xSeries, and zSeries.

IBM product engineering (PE)
The third-level of IBM service support. Product engineering is composed of IBM engineers who have experience in supporting a product or who are knowledgeable about the product.

IBM Serial Storage adapter
A physical adapter based on the IBM Serial Storage architecture. IBM Serial Storage adapters connect disk drive modules to DS6000 clusters.

IBM TotalStorage
The brand name used to identify storage products from IBM, including the IBM TotalStorage DS6000. See also *IBM TotalStorage DS6000* and *IBM TotalStorage DS Storage Manager*.

IBM TotalStorage DS6000
A member of the IBM TotalStorage 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 TotalStorage DS CLI
The command-line interface (CLI) that is specific to the DS6000.

IBM TotalStorage 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 TotalStorage 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 TotalStorage Multipath Subsystem Device Driver (SDD)
Software that is designed to support the multipath configuration environments of the DS6000. The SDD resides in a host system with the native disk device driver.

IBM TotalStorage Resiliency Family
A set of hardware and software features and products, as well as integrated software and services that are available on theIBM TotalStorage 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 X3.230-199*x*. The IP acronym is the IP in TCP/IP. See also *Transmission Control Protocol/Internet Protocol*.

invalidate
To remove a logical data unit from cache memory because it cannot support continued access to the logical data unit on the device. This removal might be the result of a failure within the storage server or a storage device that is associated with the device.

I/O See *input/output*.

I/O adapter (IOA)
An input-output adapter on the PCI bus.

IOCDS
See *input/output configuration data set*.

IOCP See *I/O Configuration Program*.

I/O Configuration Program (IOCP)
A program that defines to a system all the available I/O devices and channel paths.

I/O device
An addressable read and write unit, such as a disk drive device, magnetic tape device, or printer.

I/O interface
An interface that enables a host to perform read and write operations with its associated peripheral devices.

I/O Priority Queueing
A facility in the Workload Manager of zSeries that enables the system administrator to set priorities for queueing I/Os from different system images. See also *multiple allegiance* and *parallel access volume*.

I/O processor (IOP)
Controls input-output adapters and other devices.

I/O sequential response time
The time an I/O request is queued in processor memory waiting for previous I/Os to the same volume to complete.

IOSQ See *I/O sequential response time*.

IP See *Internet Protocol*.

IPL See *initial program load*.

iSeries
An IBM e(logo)server product that emphasizes integration. It is the successor to the AS/400 family of servers.

J

Java Virtual Machine (JVM)
A software implementation of a central processing unit (CPU) that runs compiled Java code (applets and applications). (GC)

JVM See *Java Virtual Machine*.

K

KB See *kilobyte*.

key field
The second (optional) field of a count key data record. The key length is specified in the count field. The key length determines the field length. The program writes the data in the key field and uses the key field to identify or locate a given record. The subsystem does not use the key field.

kilobyte (KB)
1) For processor storage, real, and virtual storage, and channel volume, 2^{10} or 1024 bytes.

2) For disk storage capacity and communications volume, 1000 bytes.

Korn shell
Interactive command interpreter and a command programming language.

KPOH See *thousands of power-on hours*.

L

LAN See *local area network*.

last-in first-out (LIFO)
A queuing technique in which the next item to be retrieved is the item most recently placed in the queue. (A)

LBA See *logical block address*.

LCU See *logical control unit*.

least recently used (LRU)
1) The algorithm used to identify and make available the cache space that contains the least-recently used data.

2) A policy for a caching algorithm that chooses to remove from cache the item that has the longest elapsed time since its last access.

LED See *light-emitting diode*.

LMC See *licensed machine code*.

licensed machine code (LMC)
Microcode that IBM does not sell as part of a machine, but licenses to the customer. LMC is implemented in a part of storage that is not addressable by user programs. Some IBM products use it to implement functions as an alternate to hard-wired circuitry.

LIFO See *last-in first-out*.

light-emitting diode (LED)
A semiconductor chip that gives off visible or infrared light when activated.

link address
On an ESCON interface, the portion of a source or destination address in a frame that ESCON uses to route a frame through an ESCON director. ESCON associates the link address with a specific switch port that is on the ESCON director. Equivalently, it associates the link address with the channel subsystem or control unit link-level functions that are attached to the switch port.

link-level facility
The ESCON hardware and logical functions of a control unit or channel subsystem that allow communication over an ESCON write interface and an ESCON read interface.

local area network (LAN)
A computer network located on a user's premises within a limited geographic area.

local e-mail
An e-mail configuration option for storage servers that are connected to a host-system network that does not have a domain name system (DNS) server.

logical address
On an ESCON interface, the portion of a

source or destination address in a frame used to select a specific channel-subsystem or control-unit image.

logical block address (LBA)
The address assigned by the 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.

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 up to 16 CKD-formatted logical subsystems (4096 CKD logical devices) and also up to 16 fixed-block logical subsystems (4096 fixed-block logical devices). The logical subsystem facilitates configuration of the 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 theDS6000 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 TotalStorage Management Console*.

Management Information Base (MIB)
1) A collection of objects that can be accessed by means of a network management protocol. (GC)
2) The MIB record conforms to the Open Systems Interconnection (OSI) standard defined by the International Organization for Standardization (ISO) for the exchange of information. See also *simple network management protocol*.

MAP See *maintenance analysis procedure*.

master storage unit
The physical unit that controls the creation of consistency groups in a Global Mirror session. The master storage unit sends commands to subordinate storage units. A storage unit can be a master for only one Global Mirror session. Contrast with *subordinate storage unit*.

maximum consistency group drain time
The value in seconds that indicates the maximum time that writes from the local site are delayed to the remote site while the current consistency group is being formed at the remote site. When this time is exceeded, the current attempt to form a consistency group is ended and another attempt is started. If this time is exceeded five times, this maximum time is ignored on the next attempt to form a consistency group. The default value is the larger of four minutes or two times the consistency group interval time if this value is set to zero.

maximum coordination time
The value in milliseconds that indicates the maximum time that is allowed for host I/O to be delayed during the coordination of the primary volumes of an Global Mirror session. The default is 50 milliseconds if this value is set to zero.

MB See *megabyte*.

MC See *IBM TotalStorage Management Console*.

MCA See *Micro Channel architecture*.

MDM See *Multiple Device Manager*.

mean time between failures (MTBF)
1) A projection of the time that an individual unit remains functional. The time is based on averaging the performance, or projected performance, of a population of statistically independent units. The units operate under a set of conditions or assumptions.

2) For a stated period in the life of a functional unit, the mean value of the lengths of time between consecutive failures under stated conditions. (I) (A)

medium
For a storage unit, the disk surface on which data is stored.

megabyte (MB)
1) For processor storage, real and virtual storage, and channel volume, 2^{20} or 1 048 576 bytes.

2) For disk storage capacity and communications volume, 1 000 000 bytes.

Metro Mirror
A function of a storage server that maintains a consistent copy of a logical volume on the same storage server or on another storage server. All modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume. See also *Remote Mirror and Copy* and *Global Copy*.

MES See *miscellaneous equipment specification*.

MIB See *management information base*.

Micro Channel architecture (MCA)
The rules that define how subsystems and adapters use the Micro Channel bus in a computer. The architecture defines the services that each subsystem can or must provide.

Microsoft Internet Explorer (MSIE)
Web browser software manufactured by Microsoft.

migration
The replacement of a system or subsystem with a different type of system or subsystem, such as replacing a SCSI host adapter with a fibre-channel host adapter. In the context of data migration regarding the 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*.

MSIE See *Microsoft Internet Explorer*.

MTBF See *mean time between failures*.

Multipath Subsystem Device Driver
See *IBM TotalStorage 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 TotalStorage Management Console that has multiple ports to which aDS6000 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.

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 anotherDS6000.

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 EsconNet pseudo host icon represents the ESCON protocol. The pseudo host icon labelled `Anonymous` represents hosts connected through the FCP protocol. *Anonymous host* is a commonly used synonym for *pseudo host*. The 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.

Single-Byte Command Code Sets Connection (SBCON)
The ANSI standard for the ESCON I/O interface.

small computer system interface (SCSI)
A standard hardware interface that enables a variety of peripheral devices to communicate with one another. (GC)

smart relay host
A mail relay or mail gateway that has the capability to correct e-mail addressing problems.

SMIT See *System Management Interface Tool*.

SMP See *symmetrical multiprocessor*.

SNMP See *Simple Network Management Protocol*.

SNMP agent
A server process that resides on a network node and is responsible for communicating with managers regarding that node. The node is represented as a managed object, which has various fields or variables that are defined in the appropriate MIB.

SNMP manager
A managing system that runs a managing application or suite of applications. These applications depend on Management Information Base (MIB) objects for information that resides on the managed system. Managers generate requests for this MIB information, and an SNMP agent on the managed system responds to these requests. A request can either be the retrieval or modification of MIB information.

software transparency
Criteria applied to a processing environment that states that changes do not require modifications to the host software in order to continue to provide an existing function.

source device
One of the devices in a dual-copy or remote-copy volume pair. All channel commands to the logical volume are directed to the source device. The data on the source device is duplicated on the target device. See also *target device.*

spare A disk drive on the 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 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. TheDS6000 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

TAP See *Telocator Alphanumeric Protocol*.

target A SCSI device that acts as a subordinate to an initiator and consists of a set of one or more logical units, each with an assigned logical unit number (LUN). The logical units on the target are typically I/O devices. A SCSI target is analogous to a zSeries control unit. See also *small computer system interface*.

target device
One of the devices in a dual-copy or remote-copy volume pair that contains a duplicate of the data that is on the source device. Unlike the source device, the target device might only accept a limited subset of data. See also *source device.*

TB See *terabyte*.

TCP/IP
See *Transmission Control Protocol/Internet Protocol*.

Telocator Alphanumeric Protocol (TAP)
An industry standard protocol for the input of paging requests.

terabyte (TB)
1) Nominally, 1 000 000 000 000 bytes, which is accurate when speaking of bandwidth and disk storage capacity.

2) For DS6000 cache memory, processor storage, real and virtual storage, a terabyte refers to 2^{40} or 1 099 511 627 776 bytes.

terminal emulator
A function of the management console that allows it to emulate a terminal.

thousands of power-on hours (KPOH)
A unit of time used to measure the mean time between failures (MTBF).

time sharing option (TSO)
An operating system option that provides interactive time sharing from remote terminals.

TotalStorage
See *IBM TotalStorage*.

TPF See *transaction processing facility*.

track A unit of storage on a CKD device that can be formatted to contain a number of data records. See also *home address*, *track-descriptor record*, and *data record*.

track-descriptor record (R0)
A special record on a track that follows the home address. The control program uses it to maintain certain information about the track. The record has a count field with a key length of zero, a data length of 8, and a record number of 0. This record is sometimes referred to as R0.

transaction processing facility (TPF)
A high-availability, high-performance IBM operating system, designed to support real-time, transaction-driven applications. The specialized architecture of TPF is intended to optimize system efficiency, reliability, and responsiveness for data communication and database processing. TPF provides real-time inquiry and updates to a large, centralized database, where message length is relatively short in both directions, and response time is generally less than three seconds. Formerly known as the Airline Control Program/Transaction Processing Facility (ACP/TPF).

Transmission Control Protocol (TCP)
A communications protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for internetwork protocol. TCP provides a reliable host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

Transmission Control Protocol/Internet Protocol (TCP/IP)
1) A combination of data-transmission protocols that provide end-to-end connections between applications over interconnected networks of different types.

2) A suite of transport and application protocols that run over the Internet Protocol. (GC) See also *Internet Protocol* and *Transmission Control Protocol*.

transparency
See *software transparency*.

TSO See *time sharing option*.

turbo processor
A faster multiprocessor that has six processors with common access to the main storage.

U

UFS UNIX filing system.

Ultra-SCSI
An enhanced small computer system interface.

unconfigure
To delete the configuration.

unit address
For zSeries, the address associated with a device on a given control unit. On ESCON interfaces, the unit address is the same as the device address. On OEMI interfaces, the unit address specifies a control unit and device pair on the interface.

unprotected volume
An AS/400 term that indicates that the AS/400 host recognizes the volume as an unprotected device, even though the storage resides on a RAID-formatted array and is, therefore, fault tolerant by definition. The data in an unprotected volume can be mirrored. Also referred to as an *unprotected device*.

upper-layer protocol
The layer of the Internet Protocol (IP) that supports one or more logical protocols (for example, a SCSI-command protocol and an ESA/390 command protocol). Refer to ANSI X3.230-199*x.*

UTC See *Coordinated Universal Time*.

V

virtual machine facility
A virtual data processing machine that appears to the user to be for the exclusive use of that user, but whose functions are accomplished by sharing the resources of a shared data processing system. An alternate name for the VM/370 IBM operating system.

vital product data (VPD)
Information that uniquely defines the system, hardware, software, and microcode elements of a processing system.

VM The root name of several IBM operating systems, such as VM/XA, VM/ESA, VM/CMS, and z/VM. See also *virtual machine facility*.

volume
For zSeries, the information recorded on a single unit of recording medium. Indirectly, it can refer to the unit of recording medium itself. On a nonremovable-medium storage device, the term can also indirectly refer to the storage device associated with the volume. When multiple volumes are stored on a single storage medium transparently to the program, the volumes can be referred to as logical volumes.

volume group
A collection of either physical or logical volumes.

volume label
A unique identifier that a user assigns to a logical volume.

VPD See *vital product data*.

VSE/ESA
An IBM operating system, the letters of which represent virtual storage extended/enterprise systems architecture.

W

weight distribution area
The area that is required to distribute the weight of the storage unit.

worldwide node name (WWNN)
A unique 64-bit identifier for a host that contains a fibre-channel port. See also *worldwide port name*.

worldwide port name (WWPN)
A unique 64-bit identifier associated with a fibre-channel adapter port. It is assigned in an implementation- and protocol-independent manner. See also *worldwide node name*

write hit
A write operation in which the requested data is in the cache.

write penalty
The performance impact of a classical RAID-5 write operation.

WWNN
See *worldwide node name*.

WWPN
See *worldwide port name*.

X

xSeries
The product name of an IBM e(logo)server product that emphasizes industry-standard server scalability and self-managing server technologies. It is the successor to the Netfinity family of servers.

Z

z/Architecture
An IBM architecture for mainframe computers and peripherals. The IBM e(logo)server zSeries family of servers uses the z/Architecture architecture. It is the successor to the S/390 and 9672 family of servers. See also *iSeries*.

z/OS An operating system for the IBM e(logo)server product line that supports 64-bit real storage.

z/OS Global Mirror
A function of a storage server that assists a control program to maintain a consistent copy of a logical volume on another storage unit. All modifications of the primary logical volume by any attached host are presented in order to a single host. The host then makes these modifications on the secondary logical volume. This function was formerly called extended remote copy or XRC.

zSeries
An IBM e(logo)server family of servers that emphasizes near-zero downtime. IBM enterprise servers based on z/Architecture.

zSeries storage
Storage arrays and logical volumes that are defined in the DS6000 as connected to zSeries servers.

Index

A

C

D

E

F

G

H

I

K

L

M

N

O

W

Z

Readers' comments — we would like to hear from you

IBM TotalStorage DS6000
Installation, Troubleshooting, and Recovery Guide

Publication No. GC26-7678-05

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name _____ Address _____

Company or Organization _____

Phone No. _____

Readers' Comments — We'd Like to Hear from You
GC26-7678-05

IBM

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM
Information Development
Department 61D
9032 South Rita Road
Tucson, Arizona 85775-4401

Fold and Tape | **Please do not staple** | Fold and Tape

GC26-7678-05

Cut or Fold
Along Line

Part Number: 22R5934

Printed in USA

(1P) P/N: 22R5934

GC26-7678-05

Spine information:

IBM TotalStorage DS6000 Installation, Troubleshooting, and Recovery Guide