

IBM TotalStorage DS6000



Overview

IBM TotalStorage DS6000



Overview

Note:

Before using this information and the product it supports, read the information in "Notices" on page 117.

Twenty-fourth Edition (January 2007)

© Copyright International Business Machines Corporation 2004, 2007. All rights reserved.

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

Contents

Tables v

Notices and publication information vii

Safety notices vii

Environmental notices. vii

Product recycling and disposal. vii

Battery return program viii

How to send your comments ix

Chapter 1. Overview 1

Chapter 2. Introduction to the IBM System Storage DS6000 series. 3

Overview of the DS6000 series models. 4

DS6800 (Models 1750-511 and 1750-522) 4

| DS6000 expansion enclosure (Models 1750-EX1
| and 1750-EX2) 5

IBM System Storage Management Console 6

Determine storage configuration. 7

Chapter 3. Performance features 9

Chapter 4. Arrays 11

Chapter 5. Data availability features 13

RAID implementation 13

RAID 5 overview 13

RAID 10 overview 13

Licensed functions 14

Chapter 6. Subsystem device driver for open-systems. 15

Chapter 7. Balancing the I/O load 17

Chapter 8. Concurrent code load 19

Chapter 9. System Storage Resiliency Family 21

Copy Services functions across a 2105 and 1750 21

| Copy Services functions with i5/OS 22

FlashCopy. 24

Reverse FlashCopy relationships 25

FlashCopy to Metro Mirror source volumes. 26

Multiple relationship FlashCopy 27

Persistent FlashCopy 27

Resynchronization of FlashCopy volume pairs. 28

Path connections 28

Remote mirror and copy 29

Global Copy 30

Metro Mirror 31

Global Mirror. 33

Chapter 10. Host system attachment 39

FlashCopy and Remote Mirror and Copy attachment restrictions. 39

Host systems that DS6000 series supports 39

Fibre-channel host attachments 40

Open-systems hosts with fibre-channel adapters 40

FICON-attached zSeries hosts 41

Chapter 11. DS6000 hardware resources 43

Disk drive modules. 43

Front display panel. 45

Rear operator panel 47

Power supply 50

Processor cards 52

Battery backup unit. 56

System service card. 58

Chapter 12. DS6000 resource location codes 59

Chapter 13. DS6000 Interfaces 63

IBM System Storage DS Storage Manager 63

DS6000 Storage Manager navigation overview. 63

DS open application programming interface 64

The DS command-line interface. 65

| IBM TotalStorage Productivity Center for Disk. 65

IBM TotalStorage Productivity Center for

Replication 66

Chapter 14. Software requirements 69

Chapter 15. Configuring your DS6000 series 71

Logical configuration overview. 71

Simulated configuration overview. 73

Real-time configuration overview 73

Express configuration overview. 73

Chapter 16. Licensing on the DS6000 75

Chapter 17. Light path overview 77

| Chapter 18. Logical subsystems. 81

Chapter 19. Logical volume manager 83

Chapter 20. LUN calculation 85

Chapter 21. MTMS labels 89

Chapter 22. Notification methods 91

Chapter 23. Meeting your workload needs on demand	93	Logical control unit	112
Chapter 24. Open systems	95	Logical volumes	112
Chapter 25. Power control modes	97	Parallel access volumes	112
Chapter 26. Problem management	99	Base and alias exposures	113
Chapter 27. Ranks	101	Multiple allegiance for FICON hosts.	113
Chapter 28. Remote support	103	Chapter 33. Virtual Private Network	115
Chapter 29. SNMP alerts.	105	Notices	117
Chapter 30. Storage consolidation	107	Accessibility	118
Chapter 31. Volume groups	109	Trademarks	119
Chapter 32. Volume storage	111	Terms and conditions.	120
Count key data	111	Electronic emission notices	120
Extent pools	111	Federal Communications Commission (FCC)	
Fixed block	112	statement.	121
		Industry Canada compliance statement.	121
		European community compliance statement	121
		Japanese Voluntary Control Council for	
		Interference (VCCI) class A statement	122
		Korean Ministry of Information and	
		Communication (MIC) statement	122
		Taiwan class A compliance statement	122
		Index	123

Tables

1. DDM Indicators	44	8. Disk controller SFP port indicators	55
2. Indicator states for the front display panel	46	9. Storage disk expansion SFP port indicators	56
3. Indicator states for the rear operator panel	48	10. Ethernet port indicators	56
4. Rear operator panel switches and buttons	50	11. LED indicator states for the battery backup unit	58
5. Indicator descriptions for the power supply	51	12. LED descriptions.	77
6. Server enclosure and expansion enclosure processor card indicators	53	13. Capacity and models of disk volumes for IBM iSeries	87
7. Host SFP port indicators	54		

Notices and publication information

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

Safety notices

Complete this task to find information about safety notices.

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

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

DANGER

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

1000

CAUTION:

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

1001

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


Environmental notices

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

Product recycling and disposal

This unit contains recyclable materials.

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



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

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


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

Battery return program

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

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

In the Netherlands the following applies:



For Taiwan:



Please recycle batteries.

廢電池請回收

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

Chapter 1. Overview

The topics in this section provide overview information related to your DS6000. The broad range of topics covered includes accessibility, copy services, device storage, host system attachments, concurrent code loads, input/output configuration programs, and volume storage.

Chapter 2. Introduction to the IBM System Storage DS6000 series

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

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

The DS6000 series offers high scalability and excellent performance. With the DS6800 (Models 1750-511 and 1750-522), you can install up to 16 disk drive modules (DDMs). The minimum storage capability with 4 DDMs is 292 GB. The DS6000 series also offers 500 GB, 7 200 rpm Fibre-channel ATA (FATA) disk drives, so the maximum storage capability with 16 DDMs for the DS6800 model is 8 TB. FATA drives combine a fibre-channel interface with an ATA drive, which provides additional capacity to address specific application and business requirements. You can add FATA disk drives to the DS6000 series to satisfy the needs of various fixed-content, data-archival, reference data, and near-line applications that require large amounts of storage capacity at a lower cost.

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

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

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

The DS6000 series offers the following major features:

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

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

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

Overview of the DS6000 series models

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

DS6800 (Models 1750-511 and 1750-522)

The DS6000 offers the following features:

- Two fibre channel processor cards
- PowerPC 750GX 1 GHz processor
- 4 GB of cache
- Two battery backup units (one per each processor card)
- Two ac/dc power supplies with imbedded enclosure cooling units
- Eight 2 Gb/sec device ports
- Connectivity with the availability of two to eight fibre-channel/FICON host ports. The host ports autonegotiate to either 2 Gbps or 1 Gbps link speeds.
- Attachment to 7 DS6000 expansion enclosures.





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

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

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

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

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

The DS6000 series expansion enclosure contains the following features:

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



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

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

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

IBM System Storage Management Console

The IBM System Storage management console is the focal point for configuration, copy services management, and maintenance activities.

The management console is a user provided workstation, and is required for your DS6800. You must either provide a computer to use as your management console or you can optionally order a computer from IBM. This computer must meet a minimum set of hardware and operating system compatibility requirements. If a management console is ordered, it is placed as a separate order from the DS6800.

The management console is the workstation the DS Storage Manager software is installed on. This console is used to configure and manage the DS6000 series system.

You can also define multiple (peer) management consoles on the DS6000 primary management console for the sake of redundancy. Peer management consoles can also be removed by the Remove Peer function.

Determine storage configuration

Use this page to determine if you will be using a host-agent management method or a direct-management method.

- **Host-agent management method:** This method requires host-agent software that you install on the host server. The host-agent software allows the IBM System Storage DS6000 Storage Manager client program to manage the storage unit using the same fibre-channel connections between the host server and the storage subsystem. You must install at least one management station and a software agent host. The management station can be the host or a workstation on the Ethernet network.
- **Direct-management method:** This method uses Ethernet connections from a management station to each controller. You must install at least one management station. The management station can be the host or a workstation on the Ethernet network. Attach Ethernet cables to each management station (one pair per server enclosure). You will connect the cables to each controller later when you install the server enclosure.

If you are using the host-agent management method, refer to the *IBM System Storage DS6000 Host Attachment Guide* to read more about the following fibre-channel configurations. Determine which host fibre-channel configuration that you will be using.

- Point-to-point topology (direct connect)
- Switched-fabric topology
- Arbitrated loop topology

Chapter 3. Performance features

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

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

Integrated RAID controller technology

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

High availability

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

Industry standard fibre-channel disk drives

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

Parallel Access Volumes

Parallel Access Volumes (PAV) can provide significant performance enhancements in zSeries environments by enabling simultaneous processing for multiple I/O operations to the same logical volume, which can help to significantly reduce device queue delays. This is achieved by defining multiple addresses per volume. With Dynamic PAV, the assignment of addresses to volumes can be automatically managed to help the workload meet its performance objectives and reduce overall queuing. PAV is an optional feature for the DS6000 series and also requires the purchase of the FICON Attachment feature.

LUN and volume management

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

Addressing capabilities

The DS6000 series allows the following:

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

Simplified storage management for zSeries with z/OS

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

System management

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

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

Online Information Center

The online Information Center is an information database that provides the opportunity to quickly familiarize yourself with the major aspects of the DS6000 series and to easily recognize the topics for which you might require more information. It provides information regarding user assistance for tasks, concepts, reference, user scenarios, tutorials, and other types of user information. Because the information is all in one place rather than across multiple publications, you can access the information that you need more efficiently and effectively.

Chapter 4. Arrays

An array is an arrangement of related hard-disk-drive modules that have been assigned to a group.

A disk array is a group of disk drive modules (DDMs) that are arranged in a relationship, for example, a RAID 5 or a RAID 10 array.

Chapter 5. Data availability features

This section provides information about data availability features that are supported by DS6000 series.

The DS6000 series provides the following features:

- RAID implementation
- Copy services
- Availability support for open systems, iSeries, zSeries, and pSeries hosts
- Component redundancy to eliminate single points of hardware failure, and no single point of repair other than the enclosure.

RAID implementation

RAID implementation improves data storage reliability and performance.

Redundant array of independent disks (RAID) is a method of configuring multiple disk drives in a storage subsystem for high availability and high performance. The collection of two or more disk drives presents the image of a single disk drive to the system. In the event of a single device failure, data can be read or regenerated from the other disk drives in the array.

With RAID implementation, the DS6000 series offers fault-tolerant data storage by storing the data in different places on multiple disk drive modules (DDMs). By placing data on multiple disks, input/output operations can overlap in a balanced way to improve the basic reliability and performance of the attached storage devices.

Physical capacity for the DS6000 series can be configured as RAID 5, RAID 10, or a combination of both. RAID 5 can offer excellent performance for most applications, while RAID 10 can offer better performance for selected applications, in particular, high random write content applications in the open systems environment.

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

RAID 5 overview

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

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

RAID 10 overview

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

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

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

Licensed functions

Activate licensed features before you configure your DS6000 series.

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

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

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

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

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

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

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

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

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

Chapter 6. Subsystem device driver for open-systems

The IBM System Storage Multi-path Subsystem Device Driver (SDD) supports open-systems hosts.

The Subsystem Device Driver (SDD) resides in the host server with the native disk device driver for the storage unit. It uses redundant connections between the host server and disk storage in the DS6000 series to provide enhanced performance and data availability.

Chapter 7. Balancing the I/O load

You can generally maximize the performance of an application by spreading the I/O load across clusters, arrays, and device adapters in the storage unit.

During an attempt to balance the load within the storage unit, placement of application data is the determining factor. The following resources are the most important to balance, roughly in order of importance:

- Activity to the RAID disk groups. Use as many RAID disk groups as possible for the critical applications. Most performance bottlenecks occur because a few disks are overloaded. Spreading an application across multiple RAID disk groups ensures that as many disk drives as possible are available. This is extremely important for open-system environments where cache-hit ratios are usually low.
- Activity to the clusters. When selecting RAID disk groups for a critical application, spread them across separate clusters. Because each cluster has separate memory buses and cache memory, this maximizes the use of those resources.
- Activity to the device adapters. When selecting RAID disk groups within a cluster for a critical application, spread them across separate device adapters.
- Activity to fibre-channel or FICON ports. Use the IBM System Storage Multipath Subsystem Device Driver (SDD) or similar software for other platforms to balance I/O activity across fibre-channel ports.

Note: For information about SDD, see *IBM System Storage Multipath Subsystem Device Driver User's Guide*. This document also describes the product engineering tool, the ESSUTIL tool, which is supported in the pcmpath commands and the datapath commands.

Chapter 8. Concurrent code load

Concurrent code load is a nondisruptive procedure to install firmware or software on a storage unit. The set of code loads that can be installed within a given storage complex are model dependent.

Chapter 9. System Storage Resiliency Family

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

These products and features are available on the IBM System Storage DS6000 and IBM System Storage Enterprise Storage Server® Models 750 and 800. There are a number of advanced Copy Services functions that are part of the IBM System Storage Resiliency family and are supported across various server environments such as zSeries®, pSeries®, System i and servers from Sun and Hewlett-Packard.

The following IBM System Storage Copy Services functions are available as optional features:

- Point-in-time copy, which includes IBM System Storage FlashCopy®
- Remote mirror and copy, which includes the following functions:
 - IBM System Storage Metro Mirror (previously known as Synchronous PPRC)
 - IBM System Storage Global Copy (previously known as PPRC Extended Distance)
 - IBM System Storage Global Mirror (previously known as Asynchronous PPRC)

You can manage Copy Services functions using any of the following methods:

- IBM System Storage DS CLI (command-line interface)
- IBM System Storage DS Storage Manager (Web-based interface)
- DS API (application programming interface)
- Host I/O commands from zSeries servers
- **Failover and failback operations:** In a planned or unplanned outage, such as a disaster, you can use failover and failback operations, which use Copy Services features to reduce the time that is required to synchronize remote mirror and copy volumes after switching between sites. (A failover operation is the process of switching production to a remote site. A failback operation is the process of returning production to its original location after a disaster.)

Copy Services functions across a 2105 and 1750

Copy Services functions that are performed using either the DS Storage Manager or DS CLI are designed to interact with both the 1750 and the IBM TotalStorage® Enterprise Storage Servers (ESS) Models 750 and 800.

Most Copy Services functions that are available on the ESS 2105 are also available in the 1750 and in open systems and zSeries environments.

Consider the following guidelines:

- To perform Copy Services functions between machine types 2105 and 1750, you must configure a Copy Services domain on the DS Storage Manager or DS CLI.
- To connect to the 2105 Copy Services domain on the ESS, all interfaces that you use require an authenticated login procedure to access Copy Services functions across the storage complex. The authentication is performed by using a user

name and password that was created with the ESS Specialist. Therefore, the existing user name and password that was created with the ESS Specialist for the 2105 Copy Services domain for which you will be working must match the user name and password on the management console that is connected to the 1750. Otherwise, you must add them using either the DS Storage Manager or DS CLI as part of the procedure for adding a 2105 Copy Services domain to the storage complex.

- To manage Copy Services across the 2105 and 1750, you must install licensed internal code version 2.4.2 or later on the ESS Model 800.
- The 1750 does not support remote mirror and copy (formerly PPRC) operations with an ESCON[®] link. Therefore, if you want to configure a remote mirror and copy relationship between the 1750 and 2105, you must use a FCP link.

Copy Services functions with i5/OS

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

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

Many initial configuration tasks to use i5/OS with the DS6000 are the same across the various operating systems. However, unlike other operating systems, there are some storage requirements that are unique to i5/OS. For example, LUNs on i5/OS are defined as fixed block storage and only specific LUN sizes are supported, such as 8.5 GB, 17.5 GB, and 35.1 GB. i5/OS LUNs make available a 520-byte block to the host. i5/OS refers to volumes as LUNs.

There are two types of LUN protection: protected and unprotected. In reality, both volume protection types have the same DS6000 internal RAID protection. If you define a LUN as protected, it is protected by the hardware (the DS6000) rather than by i5/OS host-based mirroring. If you define a LUN as unprotected, either the i5/OS or another LUN of equal capacity, either internally or externally mirrors it. The unprotected type is typically used for Load Source Unit mirroring. (Each System i[™] or LPAR partition has a special disk called the Load Source Unit. It holds the system's microcode and is used to reboot the system.) Unless you intend to use i5/OS (host based) mirroring, you should define your LUNs as protected.

In addition, i5/OS supports multipathing, which is part of the base operating system. Multipathing provides redundant path solutions for System i but is only used for external storage like the DS6000. Each connection for a multipath disk unit functions independently. Several connections provide availability by allowing disk storage to be used even if a single path fails.

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

Assuming that you have set up and configured your DS6000 for i5/OS, you can perform FlashCopy and remote mirror and copy operations to create copies of

System i disk pools within a single DS6000 or to a distant DS6000 using remote mirror and copy functions. You can save this data and back it up to tape.

You can install and run the DS CLI on i5/OS to allow you to perform Copy Services functions that are supported on the DS6000. You can also use the DS Storage Manager for i5/OS to perform Copy Services functions. The size of the target volumes must be equal the size of the corresponding source volumes. Any mismatch causes the task to fail.

iSeries™ Copy Services Toolkits

The iSeries Copy Services Toolkit allows iSeries users to use the data mirroring capabilities of Copy Services functions on the DS6000. The Toolkit allows iSeries users to take advantage of these functions with minimal downtime of their production system. The Toolkit interfaces with the DS CLI and DS Storage Manager to send Copy Services commands to the DS6000. For more information about the Toolkit, contact your IBM representative.

FlashCopy with the iSeries Copy Services Toolkit

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

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

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

Remote mirror and copy with the iSeries Copy Services Toolkit

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

Backup Recovery and Media Services with FlashCopy

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

Using BRMS, you can manage your most critical and complex backups simply and easily. You can also recover your system fully in the event of a disaster or failure. In addition to these backup and recovery features, BRMS enables you to track all

of your backup media from creation to expiration. You no longer have to keep track of which backup items are on which volumes or worry that you will accidentally write over active data.

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

For more about BRMS, visit the i5/OS Information Center from the IBM Publications Web site using the following URL:

<http://publib.boulder.ibm.com/iseres/>

FlashCopy

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

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

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

- The amount of data being copied
- The number of background copy processes that are occurring
- The other activities that are occurring on the DS6000 series

FlashCopy supports the following copy options:

Consistency groups

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

Change recording

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

Establish FlashCopy on existing Metro Mirror source

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

Fast reverse

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

Inhibit writes to target

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

Multiple Relationship FlashCopy

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

Persistent FlashCopy

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

Refresh target volume

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

Reverse restore

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

Reverse FlashCopy relationships

The direction of a FlashCopy relationship can be reversed, where the volume that was previously defined as the target becomes the source for the volume that was previously defined as the source (and is now the target). The data that has changed is copied to the volume previously defined as the source.

You can reverse a FlashCopy relationship if you want to restore a source volume (Volume A) to a point in time before you performed the FlashCopy operation. In effect, you are reversing the FlashCopy operation so that it appears as though no FlashCopy operation ever happened. Keep in mind that the background copy process of a FlashCopy operation must complete before you can reverse volume A as the source and volume B as the target.

There might be certain circumstances when you might want to reverse an original FlashCopy relationship. For example, suppose you create a FlashCopy relationship between source volume A and target volume B. Data loss occurs on source volume A. To keep applications running, you can reverse the FlashCopy relationship so that volume B is copied to volume A.

Note: A fast reverse option that applies to a Global Mirror operation allows a FlashCopy relationship to be reversed without waiting for the background copy of a previous FlashCopy relationship to finish. A Global Mirror operation is based on existing Global Copy and FlashCopy operations at the target site.

Figure 1 on page 26 illustrates how a reverse restore operation works:

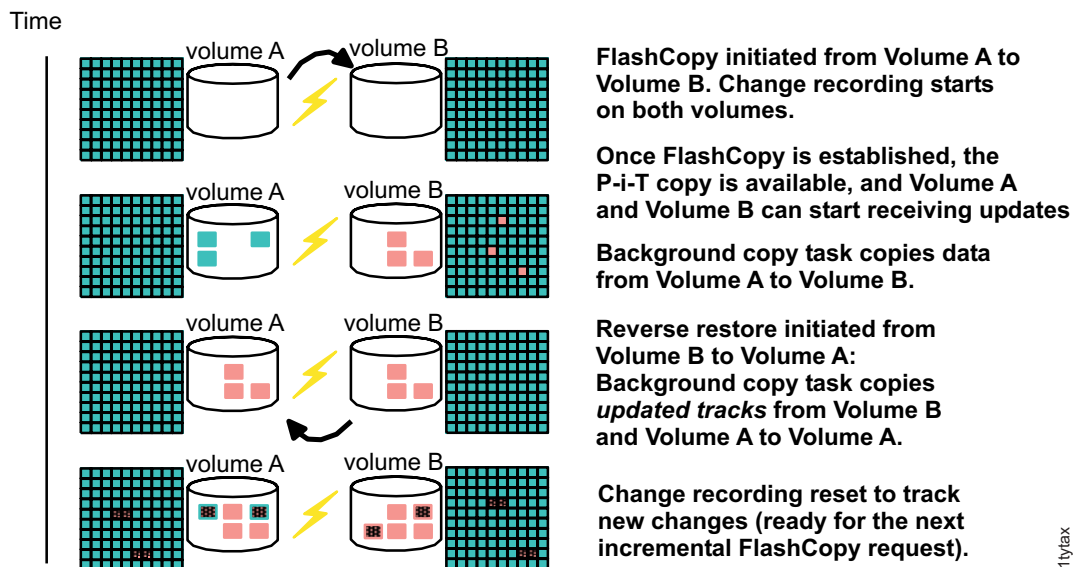


Figure 1. Refreshing target volume — reverse restore

FlashCopy to Metro Mirror source volumes

You can use an existing Metro Mirror source volume as a FlashCopy target volume. This process allows you to create a point-in-time copy and then make a copy of that data at a remote site.

Previously, you could create a FlashCopy relationship between two volumes and then create a Metro Mirror volume pair by using the FlashCopy target volume as the Metro Mirror source volume. However, you could not use an existing Metro Mirror source volume as a FlashCopy target volume. This is no longer a restriction. You can create FlashCopy target volumes that serve as production volumes and mirror them to a remote site using the Metro Mirror function.

You can create a full volume or incremental point-in-time copy to an existing Metro Mirror source volume. The first time that you create a FlashCopy operation, the entire Metro Mirror source volume has to be copied to the remote site, similar to a new pair being established. Afterwards, provided that you created the FlashCopy volume pair with the persistent and change recording options, you can perform a refresh (incremental) operation of the FlashCopy relationship. The refresh operation copies only changes and sends them to the remote site, requiring less time for the remote mirror and copy volumes to be synchronized again.

When you create a FlashCopy operation and the logical copy finishes, the corresponding Metro Mirror volume pairs turn to the duplex pending status as long as the FlashCopy data is being transferred to the Metro Mirror remote site. The time it takes to complete the copy of the FlashCopy data, until all remote mirror and copy volumes are synchronous again, depends on the amount of data to be transferred and the available bandwidth to the remote site. Be aware that, during this time, the remote site will be inconsistent if a disaster occurs. Therefore, you must carefully decide if you can tolerate not mirroring FlashCopy target volumes temporarily at a remote site.

Multiple relationship FlashCopy

The multiple relationship FlashCopy function allows a source volume to have multiple targets simultaneously.

When a FlashCopy operation copies data from a source volume to a target volume, that source volume can be involved in more than one FlashCopy relationship at the same time (known as a multiple relationship FlashCopy). If a track on a volume is not a target track of an existing FlashCopy relationship, it can become the target in a new FlashCopy relationship. Up to twelve different target FlashCopy relationships can exist for a single source volume.

If you require multiple copies of the same data, this function allows you to copy a single source volume multiple (up to twelve) times to different target volumes as soon as a FlashCopy volume is established. For example, suppose you use FlashCopy to copy volume A to volume B. As soon as that FlashCopy relationship is established, you can copy volume A to volume C. Once that relationship is established, you can copy volume A to volume D, and so on. Multiple sessions within a single volume are also possible.

Persistent FlashCopy

A persistent FlashCopy relationship is one that does not end when the background physical copy ends. The relationship persists until it is explicitly withdrawn.

Persistent FlashCopy overrides normal FlashCopy processing and allows a FlashCopy relationship to stay active until it is explicitly withdrawn. Normally, a FlashCopy relationship is automatically withdrawn when the background physical copy ends.

Designating that a FlashCopy relationship be persistent can help protect against inadvertent updates of recently created target volumes. For example, if a source volume is regularly copied to alternating target volumes (thereby ensuring that a complete copy of the source volume is always available), the persistent relationship identifies the target volume for the most recently completed FlashCopy.

Consider the following situation:

- A company wants to maintain more than one point-in-time version of a set of its production volumes.
- The company creates a FlashCopy version of its production volumes every eight hours, each time using a different set of target volumes.
- The production volume is Volume A.
- The copies are Volume B, Volume C, and Volume D.
- The volumes that are copied contain a different point-in-time version.

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

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

This FlashCopy procedure presents a possible exposure for mistakes. For instance, if the FlashCopy operation (step 2 of the FlashCopy procedure) is mistakenly

performed from Volume A to Volume B rather than from Volume A to Volume C, the result could be data that overwrites the original point-in-time copy in volume B.

The use of the persistent FlashCopy function requires that the FlashCopy relationship is explicitly withdrawn before another FlashCopy operation can be performed. While this restriction does not guarantee that a point-in-time copy cannot be overwritten by mistake, it does provide an extra measure of security.

Note: Persistent FlashCopy is always required for refreshing a FlashCopy target volume.

Resynchronization of FlashCopy volume pairs

FlashCopy volumes resynchronization (resync) provides the ability to update an initial point-in-time copy of a source volume.

You can resync your FlashCopy volumes without having to recopy your entire volume. This process is only possible with a persistent relationship, whereby the storage unit continually tracks updates to the source and target volumes. With persistent relationships, the relationship between the source and target volumes is maintained after the background copy has completed.

This function reduces the background copy completion time because only a subset of data is copied when either the source or target volume has changed. The target volume becomes current with the newly established point-in-time source copy.

To enable this function of FlashCopy, the storage unit uses a feature called change recording to monitor write operations and record changes that were made to volumes participating in FlashCopy relationships since the initial (or last) incremental FlashCopy was performed. After the initial copy of the source volume to the target volume, the FlashCopy relationship between the volumes is not automatically deleted. Rather, it remains to allow the volumes to be resynchronized. The next time that you want to update the information on your target volume, you can resync your target volume again.

Note: If you perform a resync operation, it is sent to any source and volume pair. All source and target volume pairs will be resynchronized, not just the one pair that you selected.

Path connections

Before you can create source and target remote mirror and copy pair relationships, you must establish logical paths between a logical subsystem (LSS) in a source storage unit and an LSS in a target storage unit.

When you implement a remote mirror and copy configuration, the definition of paths that remote mirror and copy uses is an important consideration. To establish paths, you must have Fibre Channel Protocol (FCP) adapters on both the source and target storage unit for the peer-to-peer links. Paths are required for communicating between volume pairs and for copying data from a source volume to a target volume.

You can establish FCP paths to connect to a source LSS and a target LSS that contain count-key-data (CKD) or fixed block (FB) volumes. The requirement is that both the source LSS and target LSS must contain the same type of volumes.

You can define from one to eight FCP paths, depending on the capability of the LSS, from a single source site LSS to a specific target site LSS.

When you establish remote mirror and copy paths over FCP links, the source storage unit uses a worldwide port name (WWPN) to communicate with the target storage unit. The worldwide node name (WWNN) and system adapters identification (SAID) of the source and destination adapters are used to generate a WWPN. A WWPN consists of exactly 16 hexadecimal digits. The hexadecimal characters are 0–9, a–f, and A–F. The values for the WWNN and the WWPN must be unique.

With FCP, a single fibre-channel link between two ports enables bi-directional remote mirror and copy operations. This means that it is possible to send and receive data in both directions simultaneously. For example, one link can have a path that is established in one direction and then have another path that is established in another direction at the same time, on the same physical path.

While a FCP interface can accommodate simultaneous data transfers, it does have limited bandwidth. To distribute workload evenly to all available paths, the storage unit monitors the overall workload on each port and selects paths that are determined by the size of the data transfer, the available bandwidth available on each FCP port, and the number of data transfers, currently processing on each port. Selecting paths in this manner ensures good response time and overall system throughput.

Remote mirror and copy

Remote mirror and copy is a storage server feature that constantly updates a target copy of a volume so that it matches changes that were made to its associated source volume.

The remote mirror and copy feature is a hardware solution that enables the mirroring of data from one site (the local site) and its associated volumes (source volumes) to a second storage unit at another site (the recovery or remote site) and its volumes (target volumes).

Note: Remote mirror and copy was referred to as Peer-to-Peer Remote Copy (PPRC) in earlier documentation for the IBM TotalStorage Enterprise Storage Server. Currently, the remote mirror and copy feature name is used generically for all methods of synchronous and asynchronous methods. In addition, in a remote mirror and copy environment, a site where production applications run might be referred to as the production site, local site, or Site A. A site that has mirrored data of your local site might be referred to as a recovery site, backup site, remote site, or Site B.

Remote mirror and copy differs from FlashCopy in two essential ways. First, the source and target volumes can be on the same storage server or on separate storage servers that are located at some distance from each other. Second, and more significantly, remote mirror and copy does not capture the state of the source volume at some point in time, but rather reflects all changes that were made on the source volume to the target volume.

Global Copy

Global Copy functions offer a nonsynchronous long-distance copy option whereby write operations to a storage unit at your production site are considered complete before they are transmitted to a storage unit at your recovery site.

Global Copy is a nonsynchronous mirroring function and is an alternative mirroring approach to Metro Mirror operations. Host updates to the source volume are not delayed by waiting for the update to be confirmed by a storage unit at your recovery site. The source volume sends a periodic, incremental copy of updated tracks to the target volume instead of a constant stream of updates. There is no guarantee that dependent write operations are transferred in the same sequence that they have been applied to the source volume. This nonsynchronous operation results in a “fuzzy copy” at the recovery site; however, through operational procedures, you can create a point-in-time consistent copy at your recovery site that is suitable for data migration, backup, and disaster recovery purposes.

To ensure that a consistent copy of the data is created, you can periodically switch from Global Copy to Metro Mirror mode. Then, either stop application I/O or freeze write applications to source volumes and wait for all pending updates to be copied to the recovery site. At this point, you can create a FlashCopy operation at your recovery site to obtain a consistent data.

The Global Copy function can operate at very long distances—well beyond the 300 km distance that is supported for Metro Mirror—and with minimal impact to applications, with the distance limited only by the network and the channel extended technology.

During a disaster, data can be restored only to the last known consistent increment that was created. This means that data that is written to the production site but is waiting to be transferred to the recovery site is lost whenever the two storage units can no longer communicate. Be aware that the use of Global Copy functions do not guarantee against data loss. Global Mirror functions, on the other hand, provide recoverable copies of production data at a long distance recovery site by continually forming consistent sets of data with no significant performance impact. This allows a quick restart at your recovery site in case of a disaster at the production site.

The following describes the Global Copy write sequence:

1. During a Global Copy operation, the storage unit at your production site captures information about updates to the source and periodically sends those updates to the target volume at the recovery site.
2. After the initial copy of tracks, the storage unit series periodically starts a synchronization cycle where all updated tracks, in ascending order from the lowest numbered track, are copied from the source volume to the target volume. The storage unit updates the target tracks with the current information for each track, regardless of the number of updates between the time that each track was last copied, the current time, and the order in which the updates occurred.
3. When this process completes, the cycle is repeated. There is little response time degradation on application write operations in extended distance mode.
4. Write updates to the source volume receive an immediate completion because the synchronization cycle is independent of the updates to the source volume.

Metro Mirror

Metro Mirror functions offer a synchronous long-distance copy option that constantly updates a secondary copy of a volume to match changes made to a source volume.

With Metro Mirror copying, the source and target volumes can be on the same storage unit or on separate storage units. You can locate the storage unit at another site some distance away. Synchronous mirroring means that each update to the source storage unit must also be updated in the target storage unit before another update can process. When Metro Mirror receives a host update to the source volume, it completes the corresponding update to the target volume. This guarantees data consistency by ensuring that a write operation that completes is received by the host application after the update has been committed to the target storage unit and acknowledged by both the source and target storage units. (Typically, the target volumes are on a different storage unit.) This results in near perfect data consistency but can result in lag time between transactions.

With Metro Mirror, consistency is guaranteed across all volumes on which an application does write operations as long as all volume pairs are in full duplex state. When error conditions affect some of the volume pairs (or different volume pairs at different time), this consistency might be lost. For example, if one of the target volumes cannot be updated because a path fails, the corresponding source volume normally goes into a suspended state, but still allows updates. However, these updates are no longer transferred to the target volume. Only the bitmap of changed tracks is created and maintained. So the consistency across volumes is lost, although the order of write operations is still guaranteed for the other target volumes.

Metro Mirror copying supports a maximum distance of 300 km (186 mi). Delays in response times for Metro Mirror are proportional to the distance between the volumes. However, 100% of the source data is available at the recovery site when the copy operation ends.

The following procedure describes how data is written for a Remote Mirror and Copy operation. A copy to the target storage unit is synchronous with the source volume's I/O operation.

1. An application requests a write I/O to the source storage unit. The write I/O is written into cache and nonvolatile storage (NVS).
2. Metro Mirror sends the write I/O to the target storage unit cache and NVS.
3. The storage unit at the recovery site signals that the write operation has completed when the updated data is in its cache and NVS.
4. When the storage unit at the production site receives notification from the target storage unit that the write operation has completed, it returns the I/O completed status to your application.

Consistency groups

Consistency groups are used for controlling error situations and for maintaining data consistency at a recovery site.

You can create paths from a source LSS to a target LSS that have the consistency group option enabled. All volume pairs in remote mirror and copy relationships from this source to target LSS (sharing the same paths) belong to the consistency group. Other volumes pairs are not affected.

With the consistency group option enabled, when an error occurs on any volume pairs or on the links that are associated with these LSS pairs, an alert is issued and I/O to all duplex remote mirror and copy volumes on LSS pairs will be queued either until a consistency group created operation is run or the consistency group timeout time expires, allowing external automation to use the consistency group created operation to create a dependent write consistent set of target volumes over any number of LSS and disk storage units.

External automation or command-line interface commands can issue a "freeze" to all LSS pairs that relate to the application, affecting the set of source volumes in the consistency group. The freeze operation causes the source volumes to suspend and their associated paths to end. Therefore, all associated updates at the recovery site are halted, resulting in a consistent data checkpoint. The automation program is not a part of the software provided with the storage unit; it must be supplied by the user. However, IBM has offerings to assist with this automation. For more information, contact your IBM storage representative.

A consistency grouping timeout function gives automation software the time to issue the automated requests to ensure consistency. The consistency grouping timeout is a value that specifies the length of time that the volume remains in long-busy (unavailable) following an error. You can change this value by accessing Copy Services functions from the Web interface (DS Storage Manager) and navigating to the Paths page and then to the LSS options page. The default value for this parameter is set to two minutes.

Note: The consistency group option is not required for Global Mirror or for Global Copy. The Global Copy target volumes at the recovery site are likely to be inconsistent. Consistency of the volume pairs can be ensured only when all updates are transmitted and copied to the recovery site while the application write operations are quiesced. Global Mirror uses a function called a session to control the formation of consistency groups. (A session is a collection of volumes across multiple storage units that are managed together to create consistent copies of data.) The formation of these consistency groups is coordinated by the "master" storage unit, which sends commands over remote mirror and copy links to its "subordinate" storage units.

The following describes the behavior of these volume pairs when an error occurs that affects a volume in a consistency group:

- Without the consistency group option enabled, the storage unit causes the volume where the error is detected to enter a suspended state, but updates to that volume are still allowed.
- With the consistency group option enabled, the volume becomes suspended and additionally enters a long-busy state. Updates are not possible to volumes from the source LSS to target LSS that are sharing the same path.

A critical mode determines the behavior of the remote mirror and copy pairs or consistency groups after a failure in communication between the source and target storage units when all the paths between a pair are lost. You can enable this mode also by accessing Copy Services functions from the DS Storage Manager and by navigating to the Paths page and then to the LSS options page.

Failover and failback operations

The failover operation is the process of switching production to a backup facility (normally your recovery site). A failback operation is the process of returning production to its original location after a disaster or a scheduled maintenance period.

There are times, both planned and unplanned, when it is necessary to suspend disk mirroring and to make use of the secondary storage unit in your configuration. As a manual process, this can be complex. However, failover and failback recovery operations are available to simplify this process and reduce the risk of error and the time it takes to switch sites and restart I/O operations.

Failover is the process of temporarily switching production to a backup facility (normally your recovery site) following a scheduled maintenance period or a disaster at your production (or local) site. A failover operation is always followed by a failback operation, which is the process of returning production to its original location. These operations use remote mirror and copy functions to reduce the time that is required to synchronize volumes after switching sites during planned or unplanned outages.

The failover and failback operations allow change recording to be enabled on the target volumes without having to communicate between the target and source storage units. This method eliminates the need to perform a full volume copy from your recovery site to the production site, which can reduce the time that is required to resume operations at your production site.

In a typical remote mirror and copy environment, processing will temporarily failover to the storage unit at your recovery site if an outage occurs at the production site. Through use of failover operations, the state of a storage unit in your target configuration changes. As a result, the storage unit is recognized as the source storage unit in the pair. Because the failover process puts the volumes into a suspended state, changes are tracked within a bitmap. Assuming that change recording is enabled, only change data is sent to the production site to synchronize the volumes, thereby reducing the time that is required to complete the failback operation.

When it is safe to return to your production site, assuming that no physical damage has occurred to the storage unit in the location, you can delete paths and create new ones from your production site to your recovery site. Then, you can create a failback recovery request to restore the storage unit as the production storage unit in the relationship.

The following considerations are for failover and failback operations:

- The failover operation does not reverse the direction of a remote mirror and copy pair. It changes a target device into a suspended source device, while leaving the source device in its current state.
- The failback recovery operation can be issued against any remote mirror and copy volume that is in a primary suspended state. The operation copies required data from the source volume to the target volume in order to resume mirroring. Failback operations are commonly used after a failover operation has been issued to restart mirroring either in the reverse direction (remote site to local site) or in the original direction (local site to remote site).

Global Mirror

Global Mirror processing provides a long-distance remote copy solution across two sites for open systems or z/OS® (or both) data using asynchronous technology. This processing is accomplished using the DS Storage Manager or the DS command-line interface (CLI).

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

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

Global Mirror is based on existing Copy Services functions: Global Copy and FlashCopy. Global Mirror operations periodically invoke a point-in-time FlashCopy at the recovery site, at regular intervals, without disrupting the I/O to the source volume, thus giving a continuously updating, nearly up-to-date data backup. Then, by grouping many volumes into a session, which is managed by the master storage unit, you can copy multiple volumes to the recovery site simultaneously while maintaining point-in-time consistency across those volumes.

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

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

To better understand how Global Mirror works, you must be familiar with the following terms:

Master

The master storage unit controls the creation of consistency groups in the Global Mirror session. The master storage unit sends commands to subordinate storage units. A storage unit is allowed to be a master for only one Global Mirror session.

Subordinate

The subordinate storage unit receives commands from a master storage unit. The subordinate storage unit is identified when a Global Mirror session is started. The subordinate storage unit forms consistency groups and performs other Global Mirror processing. A subordinate storage unit can only be controlled by one storage unit master.

Session

A session is a collection of volumes across multiple storage units that are managed together to create consistent copies of data. The session is identified with an ID that is unique across the enterprise. The ID identifies volumes that will participate in the Global Mirror consistency group. A session is opened on each LSS in the enterprise that has or might have volumes that will participate in the Global Mirror consistency group that is associated with the specific session ID.

Control path

The control path is established from the master storage unit to the subordinate storage unit when more than one storage unit participates in the Global Mirror session. If there is only one storage unit that is involved, you do not have to create a control path. The master storage unit communicates directly with its subordinate storage unit.

Summary: How Global Mirror works

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

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

Global Mirror consistency groups

Global Mirror can control the formation of consistency groups for data consistency.

Global Mirror is based on a combination of Global Copy and FlashCopy functions. To support consistent data across storage units, Global Mirror uses a function called *sessions* for forming consistent copies. A consistency group is a collection of volumes across multiple storage units that are managed together when creating consistent copies of data. The formation of these consistency groups is coordinated by the master storage unit, which sends commands over remote mirror and copy links to its subordinate storage units.

With Global Mirror functions, consistency groups can be formed many times per hour, as opposed to once or twice per day. By combining many volumes into a session, which is managed by the Global Mirror master storage unit, multiple volumes can be copied to the recovery site simultaneously while maintaining point-in-time consistency across those volumes.

The following properties control how often consistency groups are formed. You can modify these values using the Web-based interface called the IBM System Storage DS Storage Manager. Access the Global Mirror page, and then Define Properties from the summary of actions menu.

Consistency group interval time

Indicates how long (in seconds) that the storage unit waits between the formation of consistency groups. If you set this value to zero (requires sufficient bandwidth), the consistency group forms continuously, which means that a consistency group starts forming as soon as the previous consistency group has been formed.

Maximum coordination interval

Indicates the maximum time (in milliseconds) that the master storage unit communicates with its subordinate storage units to form a consistent data point. The default is 50 milliseconds.

Because host write I/O is delayed while the consistent point is being set, performance might be affected by allowing too much time for this interval. If the time set for the maximum coordination interval expires before formation of the consistent data point is complete, the consistency group fails.

Maximum time writes are inhibited to the remote site

Indicates the maximum time (in seconds) that write operations are not allowed at the recovery site before the storage unit stops forming a current consistency group. If the drain time is maintained for an extended period of time, the number of write operations that are required to transfer the data to the recovery site can become large enough to increase the time to form a consistency group.

Note: As distance increases, there is an increasing lag in the time it takes for the data to be written to the recovery site. This lag is referred to as a *drain* time.

Loss of data can also be increased in the event of a disaster. If it is set to zero, it defaults to either four minutes or twice the consistency group interval value, depending on which value is larger. The first consistency group is formed regardless of the consistency group drain time. For the remainder of the consistency groups, if the specified time expires before the data is drained to the recovery site, consistency group formation stops. After the consistency group formation has stopped five consecutive times, the timer is disabled and the next consistency group is formed regardless of the required time.

Fast restore operations

The fast reverse option allows you to reverse a FlashCopy relationship without waiting for the background copy of a previous FlashCopy to finish.

In disaster recovery situations, when you use Global Mirror, you must use volumes with consistent data at the target site. Part of the setup requires that you enable the fast restore option, a copy option that is available for FlashCopy operations. (This option applies to a Global Mirror operation, which is based on existing Global Copy and FlashCopy operations at the target site.)

Use fast restore option to reverse a FlashCopy target volume and allow consistent data to be copied back to its associated source volume, without having to wait for the background copy from the original source to the original target to complete. You can then vary the volumes online and start your applications.

Session management

A Copy Services session is a collection of volumes across multiple storage units that are managed together to create consistent copies of data. Sessions are managed for Global Mirror relationships.

The Global Mirror mode supports data consistency across multiple volumes, logical subsystems (LSSs), and storage units by using a concept called *sessions*. For disaster recovery, it is required that the volumes in a session that are necessary for recovery be consistent up to the time of failure. If data on these volumes is copied

out of sequence, then the recovery procedures that use these volumes become useless. Global Mirror uses special algorithms to provide data consistency.

Sessions are associated Global Mirror relationships and are defined with an identifier (session ID) that is unique across the enterprise. The ID identifies volumes that are related and that can participate in the Global Mirror consistency group.

You can select a set of volumes across LSSs on multiple storage units to participate in consistency groups. Logical volumes within a given LSS can be part of different sessions and different consistency groups. After you create a session, you can add logical volumes to or remove logical volumes from, a particular session. After a set of logical volumes become part of the session, a single session-based "command" can operate on all of the logical volumes in the session.

Chapter 10. Host system attachment

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

FlashCopy and Remote Mirror and Copy attachment restrictions

Some attachment restrictions apply to FlashCopy and Remote Mirror and Copy processing.

The source and target volumes should be on different host systems when you copy a source volume to a target volume with FlashCopy or Remote Mirror and Copy. The source and target volumes should be on different host systems when you require concurrent read/write access of both volumes. A copy operation with the source and target volume on the same host system creates a target volume with the same identification as the source volume. The host system sees two identical volumes.

When the copy operation creates the same identification for the target volume as for the source volumes, you are not able to distinguish one from the other. Therefore, you might not be able to access the original data.

Note: You cannot create a host target on a single Novell NetWare host system. For Novell NetWare, the target volume must be attached to a second Novell NetWare host system.

The target volume and the source volume can be on the same host system for a Remote Mirror and Copy or FlashCopy operation only under the following conditions:

- For AIX®, when the host system is using a logical volume manager (LVM) with **recreatevg** command support.
- For AIX and Sun, when the host system is *not* using a logical volume manager.
- For HP using fibre-channel, when the host is using LVM with the **vfchigid -f** command.
- For any host system, when the host system can distinguish between a source and a target volume that have the same identification.

Host systems that DS6000 series supports

The DS6000 series provides a variety of host attachments so that you can consolidate storage capacity and workloads for open-systems hosts and zSeries hosts. The storage unit can be configured for fibre-channel adapters, for support of Fibre Channel Protocol (FCP) and Fibre Connection (FICON) Protocol.

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

- Port number

- Worldwide port name (WWPN)

You can configure switch ports and hub ports that are attached to the storage unit in more than one zone. This enables multiple system adapters to share access to the storage unit fibre-channel ports. Shared access to an storage unit fibre-channel port might come from host platforms that support a combination of bus adapter types and the operating systems. For information about host systems, operating system levels, host bus adapters, cables, and fabric support that IBM supports, see the DS6000 series **Interoperability Matrix** at:

<http://www.ibm.com/servers/storage/disk/ds6000/>.

Retrieve a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www-03.ibm.com/servers/storage/support/config/hba/index.wss>.

Fibre-channel host attachments

Fibre-channel technology supports increased performance, scalability, availability, and distance for attaching storage subsystems to network servers. Fibre-channel technology supports applications that require large amounts of disk storage that is shared by two or more servers. You can use fibre-channel to connect large amounts of disk storage to a server or cluster of servers.

The DS6000 series provides a fibre-channel connection when you install a fibre-channel SFP (shortwave or longwave) in the DS6800 model.

Fibre-channel architecture provides a variety of communication protocols on the storage server. The servers that are interconnected are referred to as *nodes*. Each node has one or more ports.

A storage unit is a node in a fibre-channel network. Each port on a DS6800 fibre-channel SFP is a fibre-channel port. A host is also a node in a fibre-channel network. Each port attaches to a serial-transmission medium that provides duplex communication with the node at the other end of the medium.

There are three basic topologies supported by fibre-channel interconnection architecture:

Point-to-point

You can use the point-to-point topology to interconnect ports directly.

Switched fabric

The switched-fabric topology provides the necessary switching functions to support communication between multiple nodes. You can use a fabric to support communication between multiple nodes.

Arbitrated loop

A Fibre Channel Arbitrated Loop (FC-AL) is a ring topology where two or more ports can be interconnected. You can use the FC-AL to interconnect up to 127 hosts on a loop. An arbitrated loop may be connected to a fabric, known as a public loop. When the loop is not connected to a fabric, it is referred to as a private loop.

Open-systems hosts with fibre-channel adapters

You can attach a DS6000 series to an open-systems host with fibre-channel adapters.

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

The IBM System Storage DS6000 series supports SAN connections at 1 Gbps to 4 Gbps with 2 Gbps host bus adapters. The DS6000 series negotiates automatically and determines whether it is best to run at 1 Gbps or 2 Gbps link speed. The IBM System Storage DS6000 series detects and operates at the greatest available link speed that is shared by both sides of the system.

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

FICON-attached zSeries hosts

The DS6000 series supports FICON-attached and zSeries hosts.

The storage unit attaches to S/390® and zSeries host systems with FICON™ channels. Each storage unit fibre-channel adapter has one port. You can configure the port to operate with the FICON upper-layer protocol. When configured for FICON, the fibre-channel port supports connections to a maximum of 128 FICON hosts. On FICON, the fibre-channel adapter can operate with fabric or point-to-point topologies. With fibre-channel adapters that are configured for FICON, the DS6800 series supports the following:

- Either fabric or point-to-point topologies
- A maximum of 128 channel connections per fibre-channel port
- A maximum of 32 logical subsystems
- A maximum of 8192 logical volumes
- A maximum of 1040 volume groups

The DS6800 supports the following operating systems for zSeries hosts:

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

FICON is an optional feature on the DS6800 system, and is available with the FICON Attachment feature number 5915.

For details about models, versions of operating systems, and releases that the DS6800 supports for these host systems, see the DS6000 series **Interoperability Matrix** at: <http://www.ibm.com/servers/storage/disk/ds6000/>

Chapter 11. DS6000 hardware resources

This section contains overview information for the hardware resources contained within your storage unit.


The following hardware resources are contained within your storage unit:

- Disk drive module (DDM)
- Front display panel
- Rear operator panel
- Power supply
- Processor card
- Battery backup unit
- System service card

Disk drive modules

The disk drive modules contain the data for your storage unit.

Disk drive modules (DDMs) are magnetic nonvolatile media that are used to store and manage application and system data. The DDMs are installed in the front of the server or expansion enclosure in a horizontal orientation with the locking handle to the left. DDMs are located in rows and columns on the front of the expansion or server enclosure.



Each server or expansion enclosure must contain a minimum of 4 DDMs. You can install a maximum of 16 DDMs per server or expansion enclosure. You must add DDMs to the storage unit in groups of 4 DDMs of the same type. DDMs must be populated in rows. The first row of 4 DDMs must be placed at the top of the enclosure followed by the second row, then the third and finally, the fourth row. The positions of the DDMs in the storage unit are shown in Figure 2.



Figure 2. Close view of the disk drive modules

If you choose to configure your storage unit with fewer than 16 DDMs, you must populate the empty DDM slots with disk drive module blanks to prevent your storage unit from overheating. A drawing of the DDM blank is shown in Figure 3.




Figure 3. DDM blank

There are different types of fibre-channel arbitrated loop (FC-AL) DDMs.

Note: Server enclosures only support one DDM type. Once the server enclosure is populated with a DDM type, all remaining DDMs must be of the same type.

Each DDM contains a Ready indicator and a Check/identify indicator. The color and state descriptions of these indicators are explained in Table 1.


Table 1. DDM Indicators

Indicator name (color)	Purpose	State descriptions
Ready (green)	Indicates the status of the DDM resource.	<ul style="list-style-type: none"> • Solid: Indicates that both ports are active and that the DDM is idle. • Flashing (fast): Indicates that the DDM is processing one or more commands and that both device adapter ports are active. • Flashing (slow): Indicates that only one device adapter port is active. • Off: The DDM is not receiving power.
Check/identify (amber)	Indicates the fault or identification status of the DDM resource.	<ul style="list-style-type: none"> • Solid: Indicates one of three states: <ul style="list-style-type: none"> – Power on diagnostics are running – A fault has been detected – The device is in service mode and is ready for removal • Flashing (slow): Identifies the DDM as a failing resource. • Off: Indicates that there is no fault or no power condition.

Front display panel

The front display panel provides information and LED indicators for the storage unit.

The front display panel provides information about the status of the server or expansion enclosure. Each enclosure has a single front display panel that is located on the left of the front of the enclosure.



The front display panel contains the power on, location, information, and system alert LED indicators. It also contains fault on rear resource, fault in external enclosure, and data cache on battery indicators. Figure 4 shows the location of the indicators. Table 2 on page 46 describes the states for the front display panel indicators.



Figure 4. Front display panel indicators

Table 2. Indicator states for the front display panel








Indicator name (color)	Purpose	State descriptions
Power on (green) 	Indicates that power is being delivered to the storage unit.	<ul style="list-style-type: none"> • Solid: The power supply has been activated and is delivering dc output power to the storage unit. • Flashing: The power supplies have been deactivated. ac power is present, but dc power output is on standby power only. • Off: The storage facility is deactivated, the power supplies are deactivated, and ac power is not present.
System location (blue) 	Indicates that the server or expansion enclosure belongs to the specified storage complex.	<ul style="list-style-type: none"> • Solid: The Identify function has been activated and all attached enclosures have illuminated their respective location indicators. • Flashing: The Identify function has been activated or deactivated and the attached enclosures are currently turning their respective location indicators on or off. • Off: The system is in normal operation and the Identify function has been deactivated.
System information (amber) 	Indicates that an event has occurred for a resource in the storage unit. You must view the logs in the DS Storage Manager to obtain further information about the event.	<ul style="list-style-type: none"> • Solid: An informational event has occurred on a resource in the current enclosure or on a resource that is contained within an attached expansion enclosure. • Flashing: The system is initializing and is not yet ready to be accessed by either the management console or host systems. Initialization is complete, and the system is ready for I/O, once the flashing amber light is turned off. • Off: No events have occurred or are currently open.
System alert (amber) 	Indicates that a fault or error has been detected on the storage complex.	<ul style="list-style-type: none"> • Solid: A fault is present within the storage complex, either in the current enclosure or an attached enclosure. • Flashing: The Remind function has been activated but the fault or error has not been corrected. • Off: The system is in normal operating condition. Faults have been corrected successfully and no other faults exist.

Table 2. Indicator states for the front display panel (continued)

Indicator name (color)	Purpose	State descriptions
Fault on rear resource (amber) 	Indicates that a fault has occurred on a resource located in the rear of the enclosure.	<ul style="list-style-type: none"> • Solid: A fault is present within the enclosure. The fault can be located in the rear of the enclosure and can be repaired through the light path. • Flashing: The Remind function has been activated but the fault or error has not been corrected. • Off: There are no faulty resources located on the rear of the enclosure that can be repaired through the light path.
Fault in external enclosure (amber) 	<p>Indicates that a fault has occurred on a expansion enclosure that is attached to the server enclosure.</p> <p>This indicator is only operational for server enclosures. It is not operational for expansion enclosures.</p>	<ul style="list-style-type: none"> • Solid: A fault exists within an attached expansion enclosure. • Flashing: The Remind function has been activated but the fault or error has not been corrected. • Off: There are no faults in any of the attached expansion enclosures that can be repaired through the light path.
Data cache on battery (green) 	<p>Indicates the state of the nonvolatile cache.</p> <p>This indicator is only operational for server enclosures. It is not operational for expansion enclosures.</p>	<ul style="list-style-type: none"> • Solid: There has been an ac power loss. Any modified data that exists is stored in the cache and is being supported by the battery backup unit. • Flashing: The charge state of the battery backup unit has changed. • Off: ac power is present and there is no modified data in the cache.

Rear operator panel

The rear operator panel provides informational LED indicators as well as methods to remind and identify resources within your storage unit.

The rear operator panel provides information about the status of the server or expansion enclosure. This panel also provides physical hardware function, such as location and reminder operations. Each enclosure has a single rear operator panel that is located at the top and center of the rear of the enclosure.



The rear operator panel contains the power button, the power on, location, information, and system alert LED indicators. The Remind and Identify function buttons as well as the fault on front resource, fault in external enclosure, and data cache on battery indicators are also located on the rear operator panel. The enclosure identifier indicator, increment button, and the rack identify function

connector are located on the right of the rear operator panel. The switches and buttons on the rear operator panel are described in Table 4 on page 50. Figure 5 shows the location of the indicators described in Table 3.



Figure 5. Rear operator panel indicators

Table 3. Indicator states for the rear operator panel




Indicator name (color)	Purpose	State descriptions
Power on (green) 	Indicates that power is being delivered to the storage unit.	<ul style="list-style-type: none"> • Solid: The power supply has been activated and is delivering dc output power to the storage unit. • Flashing: The power supplies have been deactivated. ac power is present, but dc power output is on standby power only. • Off: The storage facility is deactivated, the power supplies are deactivated, and ac power is not present.
System location (blue) 	Indicates that the server or expansion enclosure belongs to the specified storage complex.	<ul style="list-style-type: none"> • Solid: The Identify function has been activated and all attached enclosures have illuminated their respective location indicators. • Flashing: The Identify function has been activated or deactivated and the attached enclosures are currently turning their respective location indicators on or off. • Off: The system is in normal operation and the Identify function has been deactivated.
System information (amber) 	Indicates that an event has occurred for a resource in the storage unit. You must view the logs in the DS Storage Manager to obtain further information about the event.	<ul style="list-style-type: none"> • Solid: An informational event has occurred on a resource in the current enclosure or on a resource that is contained within an attached expansion enclosure. • Flashing: The system is initializing and is not yet ready to be accessed by either the management console or host systems. Initialization is complete, and the system is ready for I/O, once the flashing amber light is turned off. • Off: No events have occurred or are currently open.

Table 3. Indicator states for the rear operator panel (continued)





Indicator name (color)	Purpose	State descriptions
System alert (amber) 	Indicates that a fault or error has been detected on the storage complex.	<ul style="list-style-type: none"> • Solid: A fault is present within the storage complex, either in the current enclosure or an attached enclosure. • Flashing: The Remind function has been activated but the fault or error has not been corrected. • Off: The system is in normal operating condition. Faults have been corrected successfully and no other faults exist.
Data cache on battery (green) 	<p>Indicates the state of the nonvolatile cache.</p> <p>This indicator is only operational for server enclosures. It is not operational for expansion enclosures.</p>	<ul style="list-style-type: none"> • Solid: There has been an ac power loss. Any modified data that exists is stored in the cache and is being supported by the battery backup unit. • Flashing: The charge state of the battery backup unit has changed. • Off: ac power is present and there is no modified data in the cache.
Fault on front resource (amber) 	Indicates that a fault has occurred on a resource located in the front of the enclosure.	<ul style="list-style-type: none"> • Solid: A fault is present within the enclosure. The fault can be located in the front of the enclosure and can be repaired through the light path. • Flashing: The Remind function has been activated but the fault or error has not been corrected. • Off: There are no faulty resources located on the front of the enclosure that can be repaired through the light path.
Fault in external enclosure (amber) 	<p>Indicates that a fault has occurred on a expansion enclosure that is attached to the server enclosure.</p> <p>This indicator is only operational for server enclosures. It is not operational for expansion enclosures.</p>	<ul style="list-style-type: none"> • Solid: A fault exists within an attached expansion enclosure. • Flashing: The Remind function has been activated but the fault or error has not been corrected. • Off: There are no faults in any of the attached expansion enclosures that can be repaired through the light path.


Table 4. Rear operator panel switches and buttons

Operator function	Description
Power button	<ul style="list-style-type: none"> If the storage unit is on standby power, pressing and releasing the power button activates the power-on startup sequence. If the storage unit is in a power-on state, pressing and releasing the power button activates the power-off shutdown sequence. <p>Note: Power-on and power-off operations for expansion enclosures can be managed through the server enclosure.</p>
Remind function button	Pressing and releasing this button changes the system alert or system information indicator from a solid to a flashing state. This transition can occur only when the system alert or system information indicators are in a solid state.
Identify function button	<ul style="list-style-type: none"> If the system location indicator is off, pressing and releasing this button activates the storage facility identification function. The system location indicator enters a flashing state until all of the system location indicators on the attached expansion enclosures have been illuminated. If the system location indicator is on or flashing, pressing and releasing this button deactivates the storage facility identification function. The system location indicator enters a flashing state or continue flashing until all system location indicators on the attached expansion enclosures have been turned off.
Enclosure identifier indicator (red)	<p>Displays the identification number for the expansion enclosure. The left digit displays the device adapter loop that includes the enclosure. The right digit displays the enclosure fibre-channel arbitrated loop (FC-AL) expansion enclosure address setting.</p> <p>For all enclosures that are contained inside a storage complex, the address setting is accomplished through a programming algorithm and reflected on the LED display. You do not manually set this indicator.</p>
Enclosure address increment button	Press the increment button to set the address for an enclosure.
Rack identify function connector	This connector allows you the capability to connect the enclosure to the external rack identify hardware.

Power supply

The power supplies are the main source of power for your storage unit. They also contain the fans that circulate the air to keep your storage unit cool.

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





The redundant dual ac/dc power supplies in each storage unit provide dc power to the enclosure resources. The power supplies perform the conversion from ac power input to dc power output that is required for the storage unit resources and functions. Each individual power supply has the capacity to independently power the entire server or expansion enclosure.

The ac power cord connects the power supplies to an external power source. This cord supplies ac power from the wall outlet to the power supply.

Note: If possible, plug each of the power supplies into separate ac power sources.

Each power supply contains a fan that performs cooling management functions. There is a redundant pair of enclosure cooling system fans that are imbedded within each power supply. The fans within each power supply are arranged so that no single fan failure deprives the system of enough airflow to cool the entire enclosure, provided that the enclosure airflow circulation boundaries remain intact.

The power supplies include the ac power present, dc power present, and service/fault LED indicators. Figure 6 shows the LED indicators that are located on the power supply. The indicators are explained in Table 5.



Figure 6. LED indicators on the power supply

Table 5. Indicator descriptions for the power supply

Indicator name (color)	Purpose	State descriptions
ac power present (green)	Indicates the status of the ac line cord power to the power supply.	<ul style="list-style-type: none"> • Solid: ac power is available to the power supply. • Off: ac power is not available to the power supply.
dc power present (green)	Indicates the status of the dc power being output from the power supply to the storage unit.	<ul style="list-style-type: none"> • Solid: The +12 V dc, +5 V dc and +3.3 V dc (standby) power outputs are available. • Flashing: dc standby (+3.3 V dc) power output is available. +12 V and +5 V power outputs are not available. • Off: dc power output is not available.

Table 5. Indicator descriptions for the power supply (continued)

Indicator name (color)	Purpose	State descriptions
Service/fault (amber)	Indicates fault conditions that have been detected within the power supply.	<ul style="list-style-type: none"> • Solid: A fault condition has been detected and this resource requires service. • Flashing: The resource has been identified by system management. • Off: No known faults exist. This is the default state for this indicator.

Processor cards

Processor cards perform the operations required to manage the resources and data for the expansion or server enclosure.

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.



The server enclosure and expansion enclosure processor cards provide control and management functions for the storage unit. Each processor card is capable of providing control and management for the entire server or expansion enclosure. For a dual processor card design, a single processor card can control the entire enclosure in the event of a processor card failure.

There are two types of processor cards. The server enclosure processor card, shown in Figure 7, can be installed in the server enclosure. It contains two disk expansion ports, two disk controller ports, power on and system alert LED indicators, an Ethernet connection port, one RJ-12 serial connection port and four host SFP ports.



Figure 7. Server enclosure processor card

The expansion enclosure processor card, shown in Figure 8, can be installed in the expansion enclosure. It contains two disk expansion ports, two disk controller ports, power on and system alert LED indicators, and an RJ-12 serial connection port.



Figure 8. Expansion enclosure processor card

Fiber optic cables and SFP connectors connect your storage unit with host systems and attach expansion enclosures to your server enclosure. The fiber optic SFP ports are the disk expansion, disk controller and host ports that are located on the processor controller card in the enclosure.

Note: The expansion enclosure processor card does not contain host SFP ports. These are only located on the server enclosure processor card.

Ethernet cables are used to connect the storage unit to your network.

Figure 9 and Figure 10 show the LED indicators that are included in the processor cards and their location on the cards. Table 6 describe the different states for the indicators on the processor cards.



Figure 9. Server enclosure processor card indicators



Figure 10. Expansion enclosure processor card indicators

Table 6. Server enclosure and expansion enclosure processor card indicators

Indicator name (color)	Purpose	State descriptions
Power on (green)	Indicates the state of the DC power available to the processor card.	<ul style="list-style-type: none"> Solid: DC power is being delivered to the processor card Flashing: Main DC power is not present. +3.3 volt standby power is available to the processor card. Off: DC input power is not present in any form.
Service/fault (amber)	Indicates that a fault condition has been detected within the processor card	<ul style="list-style-type: none"> Solid: A fault condition has been detected and this processor card requires service. Flashing: This processor card has been identified by system management. Off: No known faults exist. This is the default state for this indicator.

Table 6. Server enclosure and expansion enclosure processor card indicators (continued)

Indicator name (color)	Purpose	State descriptions
Host port SFPs	Indicate the state of the host port. This indicator is available only on the server enclosure processor card.	<p>The host SFP ports have the following three LED indicators:</p> <ul style="list-style-type: none"> • Host port activity indicator (green, top left) • Host port activity indicator (green, bottom left) • Service/fault indicator (amber, top right). The service/fault indicator can have the following three states: <ul style="list-style-type: none"> – Solid: A fault condition has been detected on the host SFP port. – Flashing: A fault condition has been detected and this host SFP port requires service. – Off: No known faults exist. The host SFP port is operating normally. <p>See Table 7 for descriptions of the states of these indicators.</p>
Disk controller port SFPs	Indicate the state of the disk controller port.	<p>The disk controller SFP ports have the following LED indicators:</p> <ul style="list-style-type: none"> • Device port activity indicator (green, top left) • Device port activity indicator (amber, top right) <p>See Table 8 on page 55 for descriptions of the states of these indicators.</p>
Storage disk expansion port SFPs	Indicate the state of the storage disk expansion ports.	<p>The storage disk expansion SFP port have the following LED indicators:</p> <ul style="list-style-type: none"> • Device port activity indicator (green, top left) • Device port activity indicator (amber, top right) <p>See Table 9 on page 56 for descriptions of the states of these indicators.</p>
Ethernet ports	Indicate the state of the Ethernet connection to the network. This indicator is only available on the server enclosure processor card.	<p>The Ethernet ports have the following LED indicators:</p> <ul style="list-style-type: none"> • Link status (left) • Link activity (right) <p>See Table 10 on page 56 for descriptions of the states of these indicators.</p>

Table 7. Host SFP port indicators

Host port activity (green, top left)	Host port activity (green, bottom left)	Host SFP port state
Off	Off	Wake up failure
Off	On	POST failure
Off	Flashing (slow)	Wake up failure monitor
Off	Flashing (fast)	Failures in POST

Table 7. Host SFP port indicators (continued)

Host port activity (green, top left)	Host port activity (green, bottom left)	Host SFP port state
Off	Flashing (irregular)	POST processing is in progress
On	Off	Failure while functioning
On	On	Failure while functioning
On	Flashing (slow)	Normal (Link online at X GHz)
On	Flashing (fast)	Normal (Link online at 2X GHz)
On	Flashing (irregular)	Not defined
Flashing (slow)	Off	Normal - link down
Flashing (slow)	On	Not defined
Flashing (slow)	Flashing (slow)	Offline for download
Flashing (slow)	Flashing (fast)	Restricted offline mode (waiting for restart)
Flashing (slow)	Flashing (irregular)	Restricted offline mode, test active
Flashing (fast)	Off	Debug monitor in restricted mode
Flashing (fast)	On	Not defined
Flashing (fast)	Flashing (slow)	Debug monitor in test fixture mode
Flashing (fast)	Flashing (fast)	Debug monitor in remote mode
Flashing (fast)	Flashing (irregular)	Not defined

Table 8. Disk controller SFP port indicators

Device port activity indicator (green, top left)	Device port activity indicator (amber, top right)	Disk controller SFP port state
On	Off	All of the following conditions exist: <ul style="list-style-type: none"> • The processor card is attached to the fibre-channel loop • The fibre-channel arbitrated loop is working (but not necessarily in use) • Initial machine load (IML) was successful
Off	Off	One of the following conditions exists: <ul style="list-style-type: none"> • IML was not successful • No cables are attached to the processor card • The device subsystem fibre-channel interface controller is resetting during a configuration • The fibre-channel connection is bad, either due to a faulty cable or a continuously resetting loop

Table 9. Storage disk expansion SFP port indicators

Device port activity indicator (green, top left)	Device port activity indicator (amber, top right)	Disk expansion SFP port state
Off	Off	This is normal port status when SFP is not installed.
Off	On	A fault condition with the SFP has been detected.
Off	Flashing	A situation has occurred that requires identification of this processor card by the service package.
Off	Flashing (fast)	Failure in POST
Off	Flashing (irregular)	POST processing is in progress
On	Off	Normal operation. There is no activity.
On	On	Port is bypassed. One of the two conditions exist: <ul style="list-style-type: none"> The SFP is installed, but the port is not receiving a valid signal The port is receiving an F8 failure notification from the attached device
On	Flashing	The port is manually bypassed.
Flashing	Off	Normal operation. The data is being transferred between the port and the attached device.

Table 10. Ethernet port indicators

Connection speed indicator (left)	Duplex status indicator (right)	Ethernet connection state
On	On	The Ethernet link is established at 100 Mbps and full duplex.
On	Off	The Ethernet link is established at 100 Mbps and half duplex.
Off	On	The Ethernet link is established at 10 Mbps and full duplex.
Off	Off	Either no Ethernet link is present or the link is established at 10 Mbps and half duplex.

Battery backup unit

The battery backup unit provides reserve power and data integrity function in the case of a loss of the main power supply.

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



Battery backup units (BBUs) provide protection against data loss in the event of a power failure. The storage unit includes a redundant battery backup unit for each ac/dc power supply. Each BBU is charged by the associated power supply. The

battery backup unit can provide enough backup power to preserve data integrity for up to 72 hours without power. The battery does not back up the memory if all nonvolatile data has already been moved to disk storage.

Note: The BBU must charge before initialization can occur.

In the event that input power is lost to the cluster when there is modified customer data that has not been moved to disk storage, the cluster converts to battery backup mode. In battery backup mode, the BBU supplies an alternate source of dc power to only those cluster components that are required to preserve data integrity and provide indication of the cluster power state. If input power is restored, the cluster automatically initiates the power-on activation sequences to recover the system and conserve battery power.

During BBU service, if, after removing a battery backup unit, you choose not to immediately replace the resource, you must fill the empty BBU slot with a battery backup unit blank to prevent your storage unit from overheating. Figure 11 shows a drawing of a battery backup unit blank. If you do not replace the battery backup unit, or if you insert a blank unit, the processor card associated with the missing battery remains offline and your server or storage unit operates in a mode that risks data loss in the event of a failure.




Figure 11. Battery backup unit blank

The battery backup unit contains a dc present indicator, a charging state indicator, and a service/fault indicator. The descriptions of the states for the BBU indicators are described in Table 11 on page 58. Figure 12 shows the location of these indicators.



Figure 12. Battery backup unit indicators

Table 11. LED indicator states for the battery backup unit

Indicator name (color)	Purpose	State descriptions
dc present (green)	Indicates that power is being delivered to the battery backup units.	<ul style="list-style-type: none"> • Solid: +5 Volt dc and +12 Volt dv power is present. • Off: No power is being delivered to the BBU.
Charging state (green)	Indicates the level of charge for the battery.	<ul style="list-style-type: none"> • Solid: The BBU is fully charged. • Flashing (slow): The BBU is charging and will provide more than 75% of the power required for 72 hours of operation. • Flashing (fast): The BBU is charging and will provide less than 75% of the power required for 72 hours of operation. • Off: The BBU dc output is unavailable.
Service/fault (amber)	Indicates that a fault condition has been detected within the battery backup unit.	<ul style="list-style-type: none"> • Solid: A fault condition has been detected within the BBU and this resource requires service. • Blinking: This resource has been identified by system management. • Off: No known faults exist. This is the default state for this indicator.

System service card

The system service card provides immediate information and instructions for removing and replacing resources in the server or expansion enclosure.

The system service card contains resource removal and replacement procedures, resource part numbers, safety warnings, and space for any notes you might have about the enclosure.

The system service card is a blue plastic card found in the rear of the enclosure, at the bottom of the resources.



Chapter 12. DS6000 resource location codes

You can use the resources location codes from the DS Storage Manager to identify the physical location of any hardware resource in the enclosure.

Each hardware resource in the DS6000 has a specific resource location code that logically identifies the exact physical location of the resource.

The resource location code is presented in the following format:
Uttttmmm.ppsssss-P1-L#, where *tttt* is the machine type, *mmm* is the model type, *ppsssss* is the machine serial number, *P1* is the resource point of reference, and *L#* is the location for the specific resource. For example, U1750511.12ABC3D-P1-C4 is the resource location code for the upper processor card on the server enclosure.

The following information shows the location of the specified resource, both visually and through the resource location code:

The front display panel is identified through the **E20** resource location code. For example, U1750511.12ABC3D-P1-E20 as shown in Figure 13.




Figure 13. Front display panel - resource location code E20

The disk drive modules (DDMs) are identified through the **Dn** resource location code, where *n* is the assigned location for the specific DDM. DDMs are numbered 1 through 16, starting from the top left, moving right across the row and continue from the left side of the next row. Figure 14 shows the numbered DDMs. For example, U1750511.12ABC3D-P1-D1 through U1750511.12ABC3D-P1-D16.




Figure 14. DDMs - resource location codes D1 through D16

The rear operator panel is identified through the **C1** resource location code. For example, U1750511.12ABC3D-P1-C1 as shown in Figure 15.



Figure 15. Rear operator panel - resource location code C1

The left-side power supply is identified through the **E1** resource location code. For example, U1750511.12ABC3D-E1 as shown in Figure 16.

Figure 16. Left-side power supply - resource location code E1

The right-side power supply is identified through the **E2** resource location code. For example, U1750511.12ABC3D-E2 as shown in Figure 17.




Figure 17. Right-side power supply - resource location code E2

The upper processor card is identified through the **C2** resource location code. For example, U1750511.12ABC3D-P1-C2 as shown in Figure 18.



Figure 18. Upper processor card - resource location code C2

The lower processor card is identified through the **C4** resource location code. For example, U1750511.12ABC3D-P1-C4 as shown in Figure 19.



Figure 19. Lower processor card - resource location code C4

The left-side battery backup unit is identified through the **E10** resource location code. For example, U1750511.12ABC3D-E10 as shown in Figure 20.



Figure 20. Left-side battery backup unit - resource location code E10

The right-side battery backup unit is identified through the **E11** resource location code. For example, U1750511.12ABC3D-E11 as shown in Figure 21.




Figure 21. Right-side battery backup unit - resource location code E11

The following examples show two different resource location codes:

- U1750511.12ABC3D-E1 : This shows a resource location code for the left power supply that is located on the server enclosure with the 1750 machine type, 511 model number, and 12ABC3D serial number.
- U1750EX1.45EFG6H-P1-D7 : This shows a resource location code for the seventh DDM that is located on the expansion enclosure with the 1750 machine type, EX1 model number, and 45EFG6H serial number.

Server enclosure processor card port location codes

The ports on the server enclosure processor card are identified by adding another suffix to the resource location code for the specific processor card. Figure 22 on page 61

page 61 shows the ports that are located on the server enclosure processor card and an example location code for each of those ports.

The disk expansion ports are located on the left side of the server enclosure processor card. They are identified through the **Cx-Ty** location code, where *x* is the location code for either the upper or lower processor card and *y* is the location code (from 10 to 13) for the specific disk expansion port on that card. For example, U1750511.12ABC3D-P1-C2-T11 is the second disk expansion port from the left on the upper processor card, and U1750511.12ABC3D-P1-C4-T13 is the fourth disk expansion port from the left on the lower processor card.

The Ethernet port is located to the right of the disk expansion ports on the server enclosure processor card. It is identified through the **Cx-T20** location code, where *x* is the location code for either the upper or lower processor card. For example, U1750511.12ABC3D-P1-C4-T20 is the Ethernet port on the lower processor card.

The host attachment ports are located on the right side of the server enclosure processor card. They are identified through the **Cx-Ty** location code, where *x* is the location code for either the upper or lower processor card and *y* is the location code (from 0 to 3) for the specific host attachment port on that card. For example, U1750511.12ABC3D-P1-C2-T1 is the third host attachment port from the right on the upper processor card, and U1750511.12ABC3D-P1-C4-T3 is the first host attachment port from the right on the lower processor card.

Figure 22 outlines each port on the server enclosure processor card. The following




Figure 22. Server enclosure processor card port location codes

list provides an example location code for each port in the figure, assuming that the figure shows the *upper* server enclosure processor card.

1. First disk expansion port from the left (left DISK EXP. OUT port):
U1750511.12ABC3D-P1-C2-T10
2. Second disk expansion port from the left (right DISK EXP. OUT port):
U1750511.12ABC3D-P1-C2-T11
3. Third disk expansion port from the left (left DISK CONTRL OUT port):
U1750511.12ABC3D-P1-C2-T12
4. Fourth disk expansion port from the left (right DISK CONTRL OUT port):
U1750511.12ABC3D-P1-C2-T13
5. Ethernet (network) port: U1750511.12ABC3D-P1-C2-T20
6. Fourth host attachment port from the right: U1750511.12ABC3D-P1-C2-T0
7. Third host attachment port from the right: U1750511.12ABC3D-P1-C2-T1
8. Second host attachment port from the right: U1750511.12ABC3D-P1-C2-T2
9. First host attachment port from the right: U1750511.12ABC3D-P1-C2-T3

Expansion enclosure processor card port location codes

The ports on the expansion enclosure processor card are identified by adding another suffix to the resource location code for the specific processor card.

Figure 23 shows the ports that are located on the expansion enclosure processor card and an example location code for each of those ports.

The disk expansion ports are located on the left side of the expansion enclosure processor card. They are identified through the **Cx-Ty** location code, where *x* is the location code for either the upper or lower processor card and *y* is the location code (from 10 to 13) for the specific disk expansion port on that card. For example, U1750EX1.12ABC3D-P1-C2-T11 is the second disk expansion port from the left on the upper processor card, and U1750EX1.12ABC3D-P1-C4-T13 is the fourth disk expansion port from the left on the lower processor card.

Figure 23 outlines each port on the expansion enclosure processor card. The



Figure 23. Expansion enclosure processor card port location codes

following list provides an example location code for each port in the figure, assuming that the figure shows the *upper* expansion enclosure processor card.

1. First disk expansion port from the left (left IN port): U1750EX1.12ABC3D-P1-C2-T10
2. Second disk expansion port from the left (right IN port): U1750EX1.12ABC3D-P1-C2-T11
3. Third disk expansion port from the left (left OUT port): U1750EX1.12ABC3D-P1-C2-T12
4. Fourth disk expansion port from the left (right OUT port): U1750EX1.12ABC3D-P1-C2-T13

Chapter 13. DS6000 Interfaces

This section describes the following interfaces:

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

IBM System Storage DS Storage Manager

The IBM System Storage DS Storage Manager is an interface that is used to perform logical configurations, service, copy services management, and for firmware upgrades.

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

You can install the DS Storage Manager as either a GUI or in unattended (silent) mode for the Windows operating systems. The unattended (silent) mode installation option allows you to run the installation program unattended. Use this method of installation to customize a response file and issue a command from a command prompt window. The response file is a template on the IBM System Storage Storage Manager CD.

The DS Storage Manager can be accessed from any location that has network access using a Web browser. It offers you the following choices that are available with the applicable license management code:

Simulated configuration

This feature allows you to create and save logical configurations and apply them to an online DS6000 series. This simulated configuration tool is installed on your server and can be used for the configuration of a DS6000 series at initial installation or for reconfiguration activities.

Real-time configuration

This feature provides you with real-time configuration support.

Express configuration

This feature provides the simplest and fastest configuration method.

Copy services

This feature allows you to process copy services functions.

DS6000 Storage Manager navigation overview

From the DS Storage Manager, you can manage your configuration files, monitor your system, manage your hardware, configure your storage, and manage Copy Services.

Under **Manage configuration files**, you can manage your simulated configuration files.

Under **Monitor system**, you can:

- review system summaries
- review physical summaries
- review long running task summaries
- review properties
- review logs
- find IBM contact information
- create and modify user accounts

Under **Manage hardware**, you can manage:

- storage complexes
- storage units
- host systems

Under **Configure storage**, you can configure:

- arrays
- ranks
- extent pools
- open-systems volumes
- zSeries volumes

Under **Copy Services**, you can manage:

- FlashCopy
- Paths
- Metro Mirror
- Global Mirror

DS open application programming interface

The IBM System Storage DS Open Application Programming Interface (API) is a nonproprietary storage management client application that supports routine LUN management activities, such as LUN creation, mapping and masking, and the creation or deletion of RAID 5 and RAID 10 volume spaces. The DS Open API supports these activities through the use of the Storage Management Initiative Specification (SMI-S), as defined by the Storage Networking Industry Association (SNIA).

The DS Open API helps integrate configuration management support into storage resource management (SRM) applications, which allow customers to benefit from existing SRM applications and infrastructures. The DS Open API also enables the automation of configuration management through customer-written applications. Either way, the DS Open API presents another option for managing storage units by complementing the use of the IBM System Storage DS Storage Manager web-based interface and the DS command-line interface.

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

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

The DS command-line interface

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

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

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

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

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

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

IBM TotalStorage Productivity Center for Disk

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

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

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

Disk Manager can do the following performance functions:

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

Through the use of data collection, threshold settings, and performance reports, performance can be monitored for the ESS, DS4000, DS6000, and DS8000, SAN Volume Controller, and other storage subsystems that support the SMI-S block service performance subprofile. The performance functions start with the data collection task, which captures performance statistics for the devices and stores the data in the database.

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

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

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

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

IBM TotalStorage Productivity Center for Replication

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

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

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

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

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

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

Chapter 14. Software requirements

To see current information on servers, operating systems, I/O adapters, and connectivity products supported by the DS6000 series, click **Interoperability Matrix** at the following DS6000 series Web site:

<http://www.ibm.com/servers/storage/disk/ds6000/>

Chapter 15. Configuring your DS6000 series

This section provides an overview of the methods that you can use to configure a DS6000 series.

You can use one of the following three methods to configure your storage complex:

- Simulated configuration
- Real-time configuration
- Express configuration

Logical configuration overview

Before you configure your DS6000 series, it is important to understand IBM terminology for storage concepts and the storage hierarchy.

In the storage hierarchy, you begin with a disk. Logical groupings of four disks form an array site. Logical groupings of one or two array sites form an array. After you define your array storage type as CKD or fixed block, you can create a rank. A rank is divided into a number of fixed-size extents. If you work with an open-systems host, an extent is 1 GB. If you work with IBM eServer™ zSeries or S/390 systems, an extent is the size of an IBM 3390 Mod 1 disk drive.

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

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

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

A logical volume is composed of one or more extents. A volume group specifies a set of logical volumes. By identifying different volume groups for different uses or functions (for example, SCSI target, FICON control unit, remote mirror and copy secondary volumes, FlashCopy targets, and Copy Services), access to the set of logical volumes that are identified by the volume group can be controlled. Volume groups map hosts to volumes. Figure 24 on page 72 shows a graphic representation of the logical configuration sequence.

When volumes are created, you must initialize logical tracks from the host before the host is allowed read and write access to the logical tracks on the volumes. An internal volume initialization process allows quicker access to logical volumes that are used as host volumes and source volumes in Copy Services relationships, such

as FlashCopy or Remote Mirror and Copy relationships. Volumes are not available as target volumes for Copy Services relationships until the volumes have been fully initialized. After volume initialization is complete and successful, you can use the target volumes to establish Copy Services relationships. Initialization requests can fail if the initialization is performed on volumes that are participating in existing Copy Services relationships.

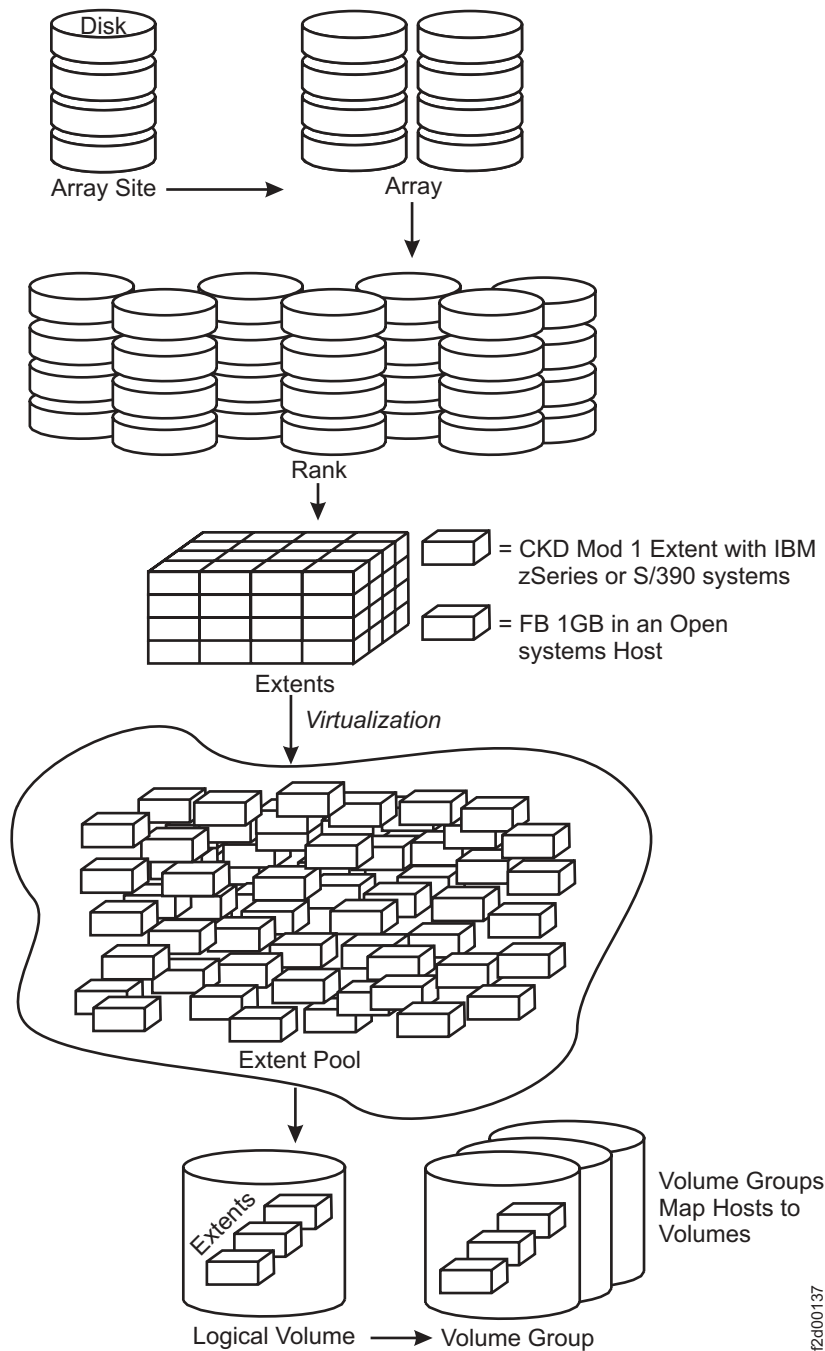


Figure 24. Logical configuration sequence

The storage management software can be used in real-time or simulated modes. When you are connected to storage devices over your network, you can use the Real-time Manager to manage your hardware or configure your storage. When you

want to work offline, you can use the Simulated Manager to create or import a configuration. This lets you export a configuration that you have already created to a system when it reaches your site.

The custom configuration option is ideal if you want to customize your configuration. The Express Configuration option shortens and simplifies the configuration process, but only allows you to create one host. This is ideal if you want only one host and you want to quickly and easily set up and configure your storage because most storage-appropriation decisions have already been made.

Simulated configuration overview

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

Use the simulated configuration method to perform the following tasks:

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

Real-time configuration overview

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

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

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

Express configuration overview

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

Some configuration methods require extensive time. Because there are many complex functions that are available to you, you are required to make several decisions during the configuration process. However, with the express

configuration method, the storage server makes several of those decisions for you, decreasing the number of steps that are required to configure volumes and hosts. This eliminates extensive configuration process time and simplifies the task for you.

The express configuration method is ideal for the following users:

- Novice users with little knowledge of storage concepts who want to quickly and easily set up and begin using storage
- Expert users who want to quickly configure a storage complex by allowing the storage server to make decisions for the best storage appropriation

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

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

Chapter 16. Licensing on the DS6000

After the IBM service representative has installed your system and before you can configure it, you must activate your licensed features on each storage unit.

The DS6000 is licensed at the following levels:

- **Machine licensing** uses licensed machine code (LMC) to license the base models and expansion models. When you receive a base or expansion model, you receive an LMC agreement. The use of the machine constitutes acceptance of the license terms outlined in the LMC agreement.
- **Operating environment licensing** manages the machine operating environment. An operating environment license (OEL) is required on every DS6000 base unit . The licensing must cover the physical capacity of the base unit and all attached expansion 4 units, excluding inactivated Standby CoD capacity.
- **Function licensing** controls the Copy Services and parallel access volume licenses. Function licensing depends on the type of data (count key data or fixed block data) that is used with the function.

You order both the operating environment and function licensing using the function authorization feature numbers.

After you purchase the function authorization features that authorize the operating environment licenses and function licenses for each storage unit, you must activate them using the following steps:

1. Obtain the activation codes from the following IBM Web site:
<http://www.ibm.com/storage/dsfa>
2. Enter the activation codes in your system.

Chapter 17. Light path overview

The light emitting diode (LED) indicators on each resource in your system provide status about informational and error events, location, and resource faults as well as other immediately required information.

The front display panel and rear operator panel provide most of the informational LEDs. The other resources contain LEDs that are specific to that resource or its function. Table 12 shows the LED icons, their meanings, the different color or flashing states each can have, and the resources that contain them.

Table 12. LED descriptions



LED icon	LED name	State descriptions	Resource containing the LED
	Power-on	A solid green light indicates that power is present on the resource. If the light is not lit, the resource is not receiving any power.	<ul style="list-style-type: none">• Front display panel• Rear operator panel• Processor card• Battery backup unit
	Location	<p>This blue light can be in a solid or blinking state. Press the Identify button on the rear operator panel of a server enclosure to identify connected enclosure enclosures. The location indicator on the server enclosure blinks as the location indicators on the connected enclosure enclosures enter a solid blue state. Once the location indicators on all of the connected enclosure enclosures are lit, the location indicator on the server enclosure becomes a solid blue light.</p> <p>If the Identify button is pressed when the server enclosure location indicator is a solid blue light, the light blinks as the location indicators on the connected enclosure enclosures are turned off. The location indicator on the server enclosure is turned off after all enclosure enclosure indicator lights are turned off.</p>	<ul style="list-style-type: none">• Front display panel• Rear operator panel

Table 12. LED descriptions (continued)







LED icon	LED name	State descriptions	Resource containing the LED
	Information	<p>A solid amber light indicates that a noncritical event has occurred. View the log in the DS Storage Manager to find more information about the event. The light remains solid until the condition has been corrected or you close the log entry in the DS Storage Manager.</p> <p>A flashing amber light indicates that the system is initializing and is not yet ready to be accessed by either the management console or host systems. Initialization is complete, and the system is ready for I/O, once the flashing amber light turns off.</p> <p>This light is lit for minor events. If the system alert indicator is lit at the same time as the information indicator, then two events (one minor and one major) have occurred on the resource.</p>	<ul style="list-style-type: none"> • Front display panel • Rear operator panel
	System alert	<p>A solid amber light indicates that a fault is present on the storage unit. Follow the other indicator LEDs to determine the location of the faulty resource. If no other resources have solid amber error LEDs, use the DS Storage Manager to view the current log entries and determine the point of failure.</p> <p>This light remains solid until the required resources have been replaced, the problem log entry has been closed through the DS Storage Manager, or unless you press the Remind button on the rear operator panel. If the Remind button is pressed on the rear operator panel, the system alert LED blinks on every 2 seconds.</p>	<ul style="list-style-type: none"> • Front display panel • Rear operator panel
	Error	<p>A solid amber light on a resource other than the front display panel or rear operator panel indicates that a fault exists on that resource. When the error LED is a solid amber, you can remove and replace the resource without any additional action on the DS Storage Manager. If a resource's error LED is lit, the system alert LEDs on the rear and front panels will also be a solid amber. The lights remain solid until the resource is replaced or the log entry is closed through the DS Storage Manager.</p>	<ul style="list-style-type: none"> • Disk drive module • Power supply • Battery backup unit • Processor card
	Data in cache on battery	<p>When this indicator is lit, there is data in the nonvolatile cache that is supported by the battery backup unit.</p>	<ul style="list-style-type: none"> • Front display panel • Rear operator panel

Table 12. LED descriptions (continued)

LED icon	LED name	State descriptions	Resource containing the LED
	Fault on opposite side	When this indicator is lit, the fault that is indicated by the system alert light is located on the rear of the enclosure. If this indicator is not lit, the fault that is indicated by the system alert light is located on the front of the enclosure.	<ul style="list-style-type: none"> • Front display panel • Rear operator panel
	Fault in external enclosure	<p>When this indicator is lit, the fault that is indicated by the system alert light is located in a connected enclosure enclosure. Use the location indicators and the Identify button on the rear operator panel to find the connected enclosure enclosure with a solid amber system alert indicator.</p> <p>If this indicator is not lit, and the system alert indicator is lit, the fault is located on the current enclosure.</p>	<ul style="list-style-type: none"> • Front display panel • Rear operator panel

The front display panel on the server enclosure contains the LEDs that display information for the entire storage unit. This includes the server enclosure as well as any connected enclosure enclosures. An LED is lit if a fault or other event occurs on any of the enclosures contained within the storage unit. Figure 25 shows the location of the LEDs on the front display panel. Table 12 on page 77 provides a description for each of the LEDs shown on the front display panel.



Figure 25. LED indicators on the front display panel

Both server enclosures and enclosure enclosures contain a rear operator panel. The rear operator panel on the server enclosure contains the LEDs that display information for the entire storage unit. The rear operator panel on the enclosure enclosure contains the LEDs that display information for the specific enclosure. An LED is lit if a fault or other event occurs on that specific enclosure. Figure 26 on page 80 shows the location of the LEDs on the rear operator panel. Table 12 on page 77

Multiple resource fault events: If there are multiple resources with a fault, and those resources can be repaired in any order, the system alert lights on the front and rear panels and the error light on each of the affected resources are a solid amber. If there are multiple resources with a fault, and those resources must be repaired in a specific order, the system alert lights on the front and rear panels are a solid amber. The error light on the resource that must be replaced first is a solid amber. Once that resource has been replaced and the error light is no longer turned on, the error light on the next resource that must be replaced is turned to a solid amber. The sequence continues in this manner until all required resources have been replaced and the system alert indicators on the front and rear panels are turned off.

Chapter 18. Logical subsystems

To facilitate configuration, a storage unit partitions its possible logical volumes into groups of volumes. Each group is referred to as a logical subsystem (LSS).

Before you can allocate storage space as logical volumes for use by specific host systems, you must define disk groups. If you are allocating storage for zSeries host systems, the format for the volumes is count key data (CKD) and for open-system environments, the format is fixed block (FB).

As part of the storage configuration process, you can configure the maximum number of LSSs that you plan to use. The DS6000 can contain up to 255 LSSs and each LSS can be connected to four other LSSs using a logical path. An LSS is a group of up to 256 logical volumes that have the same disk format, either CKD or FB. On zSeries hosts, an LSS is associated with a logical control unit image (LCU).

An LSS is uniquely identified within the storage unit by an identifier that consists two hex characters (0-9 or uppercase AF) for which the logical volumes are associated. A fully qualified LSS is designated using the storage unit identifier and the LSS identifier, such as IBM.1750-921-12FA123/1E. The LSS identifiers are important for Copy Services operations. For example, for FlashCopy operations, you specify the LSS identifier when choosing source and target volumes because the volumes can span LSSs in a storage unit.

The storage unit has a 64 KB 256 volume address space that is partitioned into 255 LSSs, where each LSS contains 256 logical volume numbers. The 255 LSS units are assigned to one of 16 address groups, where each address group contains 16 LSSs, or 4 KB volume addresses.

Storage unit functions, including some that are associated with fixed-block logical volumes, might have dependencies on LSS partitions. For example:

- The LSS partitions and their associated logical volume numbers must identify logical volumes that are specified for storage unit Copy Services operations.
- To establish Remote Mirror and Copy pairs, a logical path must be established between the associated LSS pair.
- FlashCopy pairs must reside within the same storage unit.

If you increase storage unit capacity, you can increase the number of LSSs that you have defined. This modification to increase the maximum is a nonconcurrent action. If you might need capacity increases in the future, leave the number of LSSs set to the maximum of 255.

Note: If you reduce the CKD LSS limit to zero for S/390 hosts or zSeries hosts, the storage unit does not process remote mirror and copy functions. The fixed-block LSS limit must be no lower than eight to support remote mirror and copy functions for open-systems hosts.

Chapter 19. Logical volume manager

Logical Volume Manager is a data management software primarily used with UNIX-based systems. One of its functions is to migrate data. Not all UNIX-based systems use the same LVM, however, when it comes to data migration all the versions use the same processes. The LVM can be used to migrate data at a single logical volume level, a physical volume (DDM) level, or entire contents of a volume group. The LVM software provides complete control over all disks and file systems.

The primary tasks involved in migrating data using the LVM rely upon the use of LVM commands that affect the following functions:

- Copying
- Mirroring
- Migrating

Direct copy is another method to migrate data. While this method uses LVM, its main focus is on the use of the UNIX **find** command and the **cpio** command.

The following list highlights some of the variations of LVM used by UNIX based systems.

- AIX and HP-UX 10.xx ship with logical volume management (LVM) software as part of the base operating system. LVM provides complete control over all disks and file systems that exist on an AIX system. HP-UX has similar volume management software.
- Sun Microsystems has a basic volume management product called Solstice, which is available for the Solaris systems. You can also purchase the Veritas Volume Manager (VxVM) and Veritas File System (VxFS) as optional products for Solaris.

Chapter 20. LUN calculation

The DS6000 series uses a volume capacity algorithm (calculation) to provide a logical unit number (LUN).

The following volume capacities are expressed in both binary (GB - 2^{30}) and decimal (GB - 10^9) conventions, independent of the capacity algorithm that is used to create the volume.

$$\text{GB } (2^{30}) = \text{Bytes} / 2^{30} = \text{Bytes} / 1\,073\,741\,824$$

$$\text{GB } (10^9) = \text{Bytes} / 10^9 = \text{Bytes} / 1\,000\,000\,000$$

Computer memory capacity is typically described in powers of 2; for example, MB(2^{20}), GB (2^{30}), TB (2^{40}). Computer storage capacity is typically described in powers of 10; for example, MB(10^6), GB(10^9), TB(10^{12}).

In the DS6000 family, physical storage capacities such as DDMs are generally expressed in powers of 10. Logical or effective storage capacities (logical volumes, ranks, extent pools) and processor memory capacities are expressed in powers of 2. Both of these conventions are used for logical volume effective storage capacities.

The logical volume capacity is accurate to 0.1 GB for either decimal or binary capacity convention. For logical volumes that are created with the decimal capacity algorithm, each decimal GB capacity value defines a unique volume size (the value that is used to create the volume). However, two decimal GB sizes might indicate the same binary GB size, because rounding to the nearest 0.1 GB is not sufficient to delineate all supported sizes. Similarly, two logical volumes with the same binary GB and decimal GB sizes are not necessarily the same size if they were created with a different capacity algorithm.

On open volumes with 512-byte blocks, you can determine the exact capacity of the volume in bytes by multiplying the number of blocks by 512. For open volumes with 520-byte blocks (such as iSeries volumes), you can determine the exact capacity of the volume in bytes by multiplying the number of blocks by 520. However, capacity of these volumes in bytes is expressed by multiplying the number of blocks by 512, because only 512 of the 520 bytes are used to store customer data.

On CKD volumes, the volume is created when a number of cylinders are specified. On DS6000, the basic allocation unit for CKD volumes is a CKD extent, which is exactly 1113 cylinders. The volume can be created with any number of integral cylinders in the range of 1 to 65 520 ('FFF0'X) cylinders, although volumes that are an exact multiple of 1113 cylinders do not waste any capacity in the last extent. The capacity of a CKD volume can be expressed in bytes, although the actual capacity of the volume varies depending on the format of every track on the volume. Every cylinder contains 15 tracks. A track with a 3390 track format is generally expressed as having a maximum capacity of 56 664 bytes. A track with a 3380 track format is generally expressed as having a maximum capacity of 47 476 bytes. The binary and decimal GB value that is displayed for CKD volumes is based on these capacities (for example: capacity of a volume with a 3390 track format in decimal GB = cylinders \times 15 \times 56782 / 10^9).

The logical volumes are created using one of the following three algorithms.

Binary Capacity Algorithm (B):

This algorithm provides LUNs that fully use the capacity in every extent and also are supported across the DS6000 family. The input to the GUI is the exact size of the volume in integral GB (2^{30}). The supported sizes are limited to integral multiples of the supported extent size and are further limited on larger LUNs to allow larger extent sizes in the future. The following sizes are valid:

1 to 64 GB (in 1 GB increments)

64 to 256 GB (in 4 GB increments)

256 to 1024 GB (in 16 GB increments)

1024 to 2048 GB (in 64 GB increments)

Decimal Capacity Algorithm (D):

This algorithm provides LUNs that are supported across the DS6000 family. The last extent on the volume might not be fully used because these sizes are not an integral multiple of the extent size.

You can calculate exact volume sizes as follows:

bytes = blocks \times 512 = $\text{INT} ((\text{INT} (xxx.X \times 10^9 / 512) + 63) / 64) \times 64 \times 512$, where $xxx.X$ is the value that is input by the user to the configuration process that is specified in GB (10^9). That is, this value is rounded down to a block or rounded up to a 32 KB track size.

Note: When you invoke the Metro Mirror function between the ESS 2105 and the DS6000, consider using these volumes sizes. In many applications, the primary and secondary storage unit of a Metro Mirror relationship must be exactly the same size.

The following sizes are valid:

0.1 to 982.2 GB (in 0.1 GB [10^8] increments)

iSeries Capacity Algorithm (I):

This algorithm provides iSeries LUN sizes that are specified in decimal GBs (with the assumption that each logical block contains 512 bytes of data). The set that is allowed is the complete set that is supported by iSeries.

You can define storage unit LUNs as either protected or unprotected. You can create iSeries LUNs as either logically protected or unprotected. An unprotected LUN is available for OS/400® to mirror that volume (either internal or external) to another of equal capacity. Protected logical volumes do not use OS/400, or host-based, mirroring. All IBM iSeries physical volumes are RAID 5 or RAID 10 volumes and are protected within the storage unit.

OS/400 only supports certain fixed volume sizes; for example, model sizes of 8.5 GB, 17.5 GB, and 35.1 GB. Because these volume sizes are not multiples of 1 GB, some space is left unused depending on the chosen model. iSeries LUNs make

available a 520-byte block to the host. The operating system uses 8 of these bytes; therefore, the usable space is still 512 bytes like other SCSI LUNs. The capacities that are quoted for the iSeries LUNs are normally 512-byte block storage capacity and are expressed in decimal GB (10^9). Convert these capacities to binary GB (2^{30}) to maximize the use of extents that are 1 GB (2^{30}).

Note: The supported logical volume sizes for load source units (LSUs) are 17.54 GB, 35.16 GB, 70.56 GB, and 141.1 GB. Logical volume sizes of 8.59 GB and 282.2 GB are not supported as external LSUs for iSeries.

Table 13 provides models of storage capacity and disk volumes of the IBM iSeries systems.

Table 13. Capacity and models of disk volumes for IBM iSeries

Model Number (Unprotected)	Model Number (Protected)	Capacity	Expected Number of LBAs	iSeries OS Version Support
A81	A01	8.59 GB	16 777 216 (0x01000000)	Version 5 Release 2 and Version 5 Release 3
A82	A02	17.55 GB	34 275 328 (0x020B0000)	Version 5 Release 2 and Version 5 Release 3
A85	A05	35.17 GB	68 681 728 (0x04180000)	Version 5 Release 2 and Version 5 Release 3
A84	A04	70.56 GB	137 822 208 (0x08370000)	Version 5 Release 2 and Version 5 Release 3
A86	A06	141.12 GB	275 644 416 (0x106E0000)	Version 5 Release 3 and later
A87	A07	282.25 GB	551 288 832 (0x20DC0000)	Version 5 Release 3 and later

Chapter 21. MTMS labels

The machine type, machine serial (MTMS) label assists you in identifying a specific storage unit.

The machine type, machine serial (MTMS) label contains the machine type number and the machine serial number that create a unique identification method for your server or expansion enclosure. The MTMS label is a black label with white lettering that is printed in the following format:

TYPE 1750-511
S/N 12-34567

Where 511 is the model type, 12 is the point of manufacture, and 34567 is the 5-digit serial number.

You can find the MTMS label on the flange on the front right side of the enclosure at the right front bezel.

Chapter 22. Notification methods

Call Home support and SNMP messages provide notifications to you and IBM if an event occurs on your storage unit.

Call Home

The Call Home feature allows the transmission of operational and error-related data to IBM through either the SMTP server connection or through a dial-up modem connection. The storage unit has the ability to alert IBM support to machine conditions. The DS Storage Manager sends data to IBM for error analysis and resolution.

The Call Home feature is enabled by default, but it must be configured in order to be used. You can disable the Call Home feature through the DS Storage Manager.

The following list contains the reasons or types of data for which call home support would be initiated:

- Problem notification: Data is sent when there is a problem that might require the attention of IBM support personnel.
- Heartbeat information: monitors the availability of the call home functions.
- Communication tests: to test for the successful installation and communication infrastructure.
- Informational data such as current code levels.

SNMP

Simple network management protocol (SNMP) is the standard protocol for managing networks and exchanging messages. It enables the storage unit to send external messages to notify personnel about an event.

The SNMP messages are sent from the storage unit over the local area network. You can configure the SNMP messages through the DS Storage Manager.

Chapter 23. Meeting your workload needs on demand

The DS6000 provides you with high-performance, connectivity, and reliability that allows your workload to be easily consolidated into a single storage subsystem.

The following list provides an overview of some of features that are associated with the DS6000. These features allow the DS6000 to meet and exceed your business needs.

POWER5 processor technology

The DS6000 features IBM's latest POWER5 server technology, and can provide up to six times the throughput of an ESS Model 800.

Industry standard fibre-channel disk drives

The DS6000 offers a selection of fibre-channel disk drives, including 300 GB drives, that allow a single system to scale up to 192 TB of capacity.

Four-port fibre-channel/FICON adapters

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

Host attachments

The DS6000 offers up to 32 FCP/FICON adapters, further increasing your ability to share your disk storage.

IBM Standby Capacity On Demand

Standby Capacity On Demand provides "standby" storage for the DS6000 and allows you to access the extra storage capacity whenever the need arises. With Standby CoD, IBM installs up to 64 drives in your DS6000. At any time, you can logically configure your Standby CoD capacity for use, which is a nondisruptive activity that does not require intervention from IBM. Upon logical configuration, you are charged for the capacity.

Chapter 24. Open systems

Open systems contain standardized elements that allow their use with appropriately designed applications, both local and remote.

Open systems refers, in a generic way, to the systems running Windows NT, Windows 2000, Novell NetWare, DYNIX/ptx, as well as to the systems running IBM AIX, IBM OS/400 and the many variants of non-IBM UNIX operating systems.

Chapter 25. Power control modes

The storage unit has four power on modes: standby power, dual power supply active, single power supply active, and battery backup mode.

Standby power mode

The storage unit is in standby power mode when the unit is completely installed and receiving ac power, but the power supplies are not generating dc power. Standby dc power is delivered to the power distribution system to maintain basic power management and control. This mode can occur with either one or two power supplies.

Dual power supply active mode

The storage unit is in dual power supply active mode when both power supplies are actively receiving ac power and generating both main and standby dc power. Both power supplies are sharing the power supply requirements for the entire storage unit. Dual power supply active mode is the normal operating mode for the server enclosure.

Single power supply active mode

The storage unit is in single power supply active mode when only one power supply is actively receiving ac power and generating both main and standby dc power. This single power supply is generating all of the power required to operate the storage unit.

Battery backup mode

The storage unit is in battery backup mode when ac power loss has occurred and the battery backup units are supplying dc power to only those resources that are required to maintain data integrity. When power loss occurs on a storage unit, the processor card initiates a flush of all processor hardware caches into server memory. The power-on sequence automatically begins once ac power is delivered to the power supplies.

Chapter 26. Problem management

There are a number of resources that are available to you if problems occur with your storage unit.

Problem management

You can manage fault and error notifications through any of the following methods:

- **SNMP:** This function enables external messaging when events occur on your storage unit.
- **Light path remind function:** The light path remind function allows you to set a light path reminder for resources that must be replaced.
- **Call home support:** You can enable the DS Storage Manager to send operational and problem data to IBM.

Problem determination

You can find the causes of problems through any of the following methods:

- **Problem log:** Navigate to the Logs page in the DS Storage Manager to review informational and error log entries.
- **Light path:** You can follow the light path LED indicators to locate resources that must be replaced.
- **Informational and error messages:** These messages provide additional information and actions that you can take to resolve problems or other events.

Problem analysis

You can analyze problems through any of the following methods:

- **Problem log:** Navigate to the Logs page in the DS Storage Manager to review informational and error log entries.
- **Service effects:** You can view a list of potential effects of a service action for each storage unit resource in either the Logs or status properties pages of the DS Storage Manager.
- **Light path:** You can follow the light path LED indicators to locate resources that must be replaced.
- **Remote support:** You can allow IBM support personnel to remotely access your storage unit to analyze and resolve problems.

Chapter 27. Ranks

One or more arrays are combined to create a logically contiguous storage space called a rank.

On a multi-array rank, data is striped (RAID 0) across the arrays in the rank to improve performance. The available space on each rank will be divided into extents. The extents are the building blocks of the logical volumes. An extent is striped across all disks of an array. The process of forming a rank does two things:

1. The array is defined for either fixed block (open systems) or CKD (zSeries) data. This determines the size of the set of data contained on one disk within a stripe on the array.
2. The capacity of the array is subdivided into equal sized partitions, called extents. The extent size depends on the extent type, FB or CKD.

Chapter 28. Remote support

The remote support function allows IBM Support to connect to your storage unit to assist you with problem determination, perform repairs, and correct error events.

Remote support provides a means for skilled IBM personnel who are located at a remote site to perform service on the storage unit, provided appropriate access has been enabled. It provides a communication method that enables faster problem determination and resolution.

You can enable remote support access through the DS6000 Storage Manager. Remote support uses secure shell (SSH) or the DS6000 Storage Manager to enable skilled remotely located IBM personnel to interact with your storage unit. Optionally, remote support can be done with virtual node configuration (VNC) if you download the VNC server. There are two ways to establish physical remote support to access the storage unit. The preferred way is through a connection to the local area network that is connected to the DS6000 processor card. However, IBM customer engineers can also access the storage unit through the serial port located on the rear operator panel.

Remote support is connected through a virtual private network (VPN) that you must initiate and maintain. You must initiate a VPN connection to IBM before IBM support personnel can use the remote support function. You can initiate a VPN connection through the DS6000 Storage Manager, over your local area network, or over a dial-up connection through a modem that is connected to the management console. The VPN manages any required data encryption for the data being transferred through the VPN connection.

You must grant permission for the support personnel to access your storage unit using SSH. You can use the DS6000 Storage Manager or command-line interface to enable or disable remote support over SSH. Also, you can enable or disable remote support using the secure shell or serial port. Use the guest user ID to enable or disable remote support using the secure shell or serial port.

Chapter 29. SNMP alerts

SNMP alerts inform you that a problem that has occurred on your DS6000 series.

Simple Network Management Protocol (SNMP) allows communication between SNMP managers and SNMP agents over your TCP/IP network. An SNMP manager allows you to view the messages that the SNMP agent sends. You can define your SNMP objects through the Management Information Base (MIB) in the SNMP manager. The objects that are contained in the MIB relate to the management of the storage units and attached hosts. Once these objects are defined, you can receive SNMP alerts that are sent from the SNMP agent, in this case the DS6000 series, to your SNMP manager. You receive an SNMP alert when an error occurs on your DS6000 series.

SNMP alerts are enabled after you complete the following process:

1. Install and configure an SNMP manager such as Tivoli® NetView®. You can find configuration information for your SNMP manager and MIB in the `SNMP_readme.txt` file that is located on the DS CLI installation CD-ROM.

Note: Ensure that the DS6000 processor cards and your SNMP manager can communicate through your network. The processor cards must be able to send SNMP messages to your SNMP manager. You can test your configuration using the Create test option in the DS Storage Manager or the **testcallhome** command in the DS CLI.

2. Set up your SNMP alerts in the Configure Notifications page of the DS Storage Manager, or through the **setsnmp** command of the DS CLI.

An SNMP message is sent every time a problem log entry is generated. After you receive an SNMP message, review the corresponding problem log entry in the DS Storage Manager. Follow the instructions that are provided in the problem log entry and the corresponding help documentation to resolve the problem. Most problem log entries close automatically after you correct the problem that generated the entry; however, there are a few problems that you must close manually. Problems that must be manually closed are indicated in the message text for that individual problem.

One SNMP message and one problem log entry are generated when the first instance of the problem occurs. Any additional occurrences of the same problem are noted in the problem record, but do not generate additional SNMP messages.

SNMP messages contain two types of error codes: SRCs and alert traps. A system resource code (SRC) is an 8-digit hexadecimal code. Alert traps can also use a combination of a generic alert trap number and a specific alert trap number to indicate the type of error.

Example

The following SNMP message example provides both a combination of a generic and specific alert trap and a system resource code (SRC):

```
SNMP Trap received at : 2006-01-19 10:24:13.265
Trap date is : Thu Jan 19 10:24:13 CET 2006
Message : version=0 communityString=public
enterprise=1.2.3.4.5.6.7.8.901 agentAddr=1.2.3.4 genericTrap=6
specificTrap=3 timeStamp=17968154
(1.3.6.1.4.1.2.6.130.1.5:323030362f30312f31392031313a31333a3032204d616e75666163747572657
23d49424d0a5265706f7274696e674d544d533d313735302d3531312a313341415957410a50726f624e6d3d3
23030362d30312d31392d31312e31332e30352e3733323832360a4c7061724e616d653d200a4661696c696e6
7456e636c6f737572654d544d533d313735302d3531312a313341415957410a5352433d30784245383032363
4430a4576656e74546578743d200a)
Translation of the character string after the : (one byte = ASCII character)
2006/01/19 11:13:02 Manufacturer=IBM
ReportingMTMS=1750-511*13ABCDE
ProbNm=2006-01-19-11.13.05.732826
LparName=
FailingEnclosureMTMS=1750-511*13ABCDE
SRC=0xBE80264C
EventText=
```

In this example, the SNMP message provides a generic alert trap 6 and specific alert trap 3, which indicates that a problem log entry has been created. This example SNMP message also provides an SRC of 0xBE80264C, which indicates that there is a hardware error with the host system cables or SFPs.

Chapter 30. Storage consolidation

When you use a storage unit, you can consolidate data and workloads from different kinds of independent servers into a single shared resource.

You might mix production and test servers in an open systems environment or mix S/390 or zSeries and open systems hosts. In this kind of independent server environment, servers rarely, if ever, contend for the same resource.

Although sharing resources in the storage unit has advantages for storage administration and resource sharing, there are additional implications for workload planning. The benefit of sharing is that a larger resource pool (for example, disk drives or cache) is available for critical applications. However, you must ensure that uncontrolled or unpredictable applications do not interfere with critical work. This requires the same kind of workload planning that you use when you mix various types of work on a server.

If your workload is critical, consider isolating it from other workloads. To isolate the workloads, place the data as follows:

- On separate RAID disk groups. S/390 or zSeries and open-systems data are automatically placed on separate arrays, which reduces the contention for disk use.
- On separate device adapters.
- In separate storage unit clusters, which isolates use of memory buses, microprocessors, and cache resources. Before you make this decision, verify that the isolation of your data to a single cluster provides adequate data access performance for your application.

Chapter 31. Volume groups

A volume group specifies a set of logical volumes. By identifying different volume groups for different uses (for example, SCSI target, FICON control unit, Remote Mirror and Copy secondaries, FlashCopy targets, Copy Services), access to the set of logical volumes that are identified by the volume group can be controlled.

A volume map specifies a set of logical volume numbers as a list of 2 byte logical volume numbers. The configured logical volumes have consecutive locations in the list starting from the beginning. An unconfigured entry is indicated by a X'FFFF' logical volume number. You can specify a maximum of 256 logical volumes (512 bytes). A volume mask specifies a set of logical volume numbers as a bitmap with one bit per logical volume number. A configured logical volume is indicated by a B'1'. A maximum of 65280 (63.75 KB) logical volumes can be specified (8 KB).

You can configure a CKD base address in one or more volume groups. A CKD alias address can be configured in one and only one volume group and the base address that it is associated with must be present in that volume group. The CKD alias can also be reassigned to a base address that is present within its assigned volume group. A CKD base address can be removed from a volume group if there are no alias addresses in the volume group that are associated with the base address. A CKD alias address can be removed from a volume group by deconfiguring the alias.

If an access attempt encounters an unconfigured logical volume number in the volume group, the access is treated as an access to an unconfigured logical volume. Requests to add additional volume group definitions or to add additional logical volume numbers to an existing volume group definition are rejected if the definition does not fit in the space that is allocated for volume group definitions.

Chapter 32. Volume storage

The volume storage topics cover various aspects related to managing data storage.

Count key data

In count-key-data (CKD) disk data architecture, the data field stores the user data.

Because data records can be variable in length, in CKD they all have an associated count field that indicates the user data record size. The key field enables a hardware search on a key. The commands used in the CKD architecture for managing the data and the storage devices are called channel command words.

Extent pools

An extent pool is a logical construct that is used to manage a set of extents of the same type that are associated with a given rank group.

When an extent pool is defined, it must be assigned a rank group and an extent type. One or more ranks with the same extent type can be assigned to an extent pool, but a rank can be assigned to only one extent pool. There can be as many extent pools as there are ranks. There are no constraints on the number or size of ranks assigned to the extent pool. All extents of the ranks assigned to an extent pool are independently available for allocation to logical volumes that are associated with the extent pool.

Typically, all the ranks within an extent pool have the same characteristics relative to RAID type, disk-drive module (DDM) RPM, and DDM interface rate. This allows for a differentiation of logical volumes by assigning them to the appropriate extent pool for the desired characteristics. Different extent sizes for the same device type (for example, count-key-data or fixed block) can be supported on the same storage unit, but these different extent types must be in different extent pools.

You can use the same RAID types in an extent pool. When an extent pool is defined, you must assign the following attributes:

- Server affinity
- Extent type
- RAID type

When you create extent pools, the following rules apply:

- You must configure a minimum of two extent pools to use server 0 and server 1.
- More than one rank can reside in an extent pool, but two extent pools cannot be created from only one rank. You can create one extent pool from one rank, unless the LUN capacity is greater than the capacity of one rank in the extent pool.

Review the following general considerations:

- One rank per pool will not constrain addresses.
- Ranks can be added to an extent pool at any time.
- The logical volumes defined in one extent pool can be different LSSs.

- The logical volumes in different extent pools can be in the same LSS and are only limited by the odd and even server affinity.
- Ranks can be removed from an extent pool if no extents on the rank are currently assigned to the logical volumes.
- Deleting an extent pool releases the ranks that are assigned to it and makes those ranks available for assignment to another extent pool.
- Any extent can be used to make a logical volume.
- Thresholds warn you when you are nearing the end of usable space on the extent pool.
- A Reserve Space option prevents the logical volume from being created in reserved space until the space is explicitly released.

Note: You cannot control or specify which ranks in an extent pool are used when you allocate extents to a volume.

Fixed block

In fixed block (FB) architecture, the data (the logical volumes) are mapped over fixed-size blocks or sectors.

With an FB architecture, the location of any block can be calculated to retrieve that block. This architecture uses tracks and cylinders. On a physical disk there are multiple blocks per track, and a cylinder is the group of tracks that exists under the disk heads at one point in time without performing a seek operation.

Logical control unit

A logical control unit represents a logical subsystem for zSeries hosts.

For zSeries hosts, a logical subsystem represents a logical control unit (LCU). Each control unit is associated with only one logical subsystem and groups logical volumes in groups of up to 256 logical volumes.

Logical volumes

A logical volume is the storage medium that is associated with a logical disk. It typically resides on one or more hard disk drives.

For the storage unit, the logical volumes are defined at logical configuration time. For count-key-data (CKD) servers, the logical volume size is defined by the device emulation mode and model. For fixed block (FB) hosts, you can define each FB volume (LUN) with a minimum size of a single block (512 bytes) to a maximum size of 2^{32} blocks or 2 TB.

A logical device that has nonremovable media has one and only one associated logical volume. A logical volume is composed of one or more extents. Each extent is associated with a contiguous range of addressable data units on the logical volume.

Parallel access volumes

Parallel access volumes (PAV) allow your system to access volumes in parallel when you use an S/390 or zSeries host.

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

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

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

The DS6000 series supports concurrent or parallel data transfer operations to or from the same volume from the same system or system image for zSeries hosts.

Base and alias exposures

An exposure is a device extension that enables an MVS system to perform multiple concurrent I/O operations to the device.

The storage unit uses two types of exposures:

- Base exposure: A device number that is associated with a parallel access volume (PAV). It is used for query, data control, and data access operations.
- Alias exposure: A device number that is associated with a base exposure and that represents the same logical volume as the base device number. It is used for data access I/O operations, query, and control operations.

Multiple allegiance for FICON hosts

The DS6000 series provides multiple allegiance facility support for FICON hosts.

The multiple allegiance facility enables the storage unit to accept concurrent I/O requests for a volume from multiple channel paths. This enables the storage unit to process requests from separate FICON hosts in parallel. Parallel processing of requests improves throughput and performance. The multiple allegiance facility does not require any user action.

Chapter 33. Virtual Private Network

A virtual private network (VPN) is a private network that securely connects corporate networks across the Internet to remote offices and users.

A VPN enables you to send data between two computers across a shared or public internetwork in a manner that emulates the properties of a point-to-point private link. A VPN provides user authentication, data encryption, and data integrity to ensure the security of the data while in transit across private networks and the Internet.

VPNs securely convey information across the Internet by connecting remote users, branch offices, and business partners into an extended corporate network. Many companies are replacing their existing telecommunications infrastructure with VPNs, by implementing secure IP tunnels across the Internet between corporate sites as well as to business partners and remote users.

Because security is a critical issue for companies worldwide, VPN connections provide a secure infrastructure that require systems to work together to mitigate the risk of malicious activity from both external and internal sources. Any connection from your network to the public Internet raises some of the following security concerns:

- Infection by viruses
- Intrusion by hackers
- Accessibility of your data from a remote support site
- Authorization of remote users to access your machine when a remote connection is opened

IBM VPN connections, along with the security features that are built into the DS6000 storage unit, make it possible to access IBM service representatives who can assist you in resolving complex problems without the risks that are associated with a connection to an external network. For information about the IBM VPN implementation including technical details, access the following Web site:

<http://www-1.ibm.com/support/docview.wss?rs=1114&uid=ssg1S1002693>

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.

Accessibility

Accessibility features provide users who have disabilities with the ability to successfully access information and use technology.

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features

These are the major accessibility features in the IBM System Storage DS6000 information:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. IBM Home Page Reader version 3.0 has been tested.
- You can operate features using the keyboard instead of the mouse.

Navigating by keyboard

You can use keys or key combinations to perform operations and initiate menu actions that can also be done through mouse actions. You can navigate the IBM System Storage DS6000 information from the keyboard by using the shortcut keys for your browser or Home Page Reader. See your browser Help for a list of shortcut keys that it supports. See the following Web site for a list of shortcut keys supported by Home Page Reader: http://www-306.ibm.com/able/solution_offerings/keyshort.html

Accessing the publications

You can find HTML versions of the IBM System Storage DS6000 information at the following Web site: <http://www.ehone.ibm.com/public/applications/publications/cgi-bin/pbi.cgi>

You can access the information using IBM Home Page Reader 3.0.

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
- DS4000
- DS6000
- DS8000
- e (logo)
- Enterprise Storage Server
- ES/9000
- ESCON
- FICON
- FlashCopy
- Graphically Dispersed Parallel Sysplex
- HACMP
- i5/OS
- IBM
- IntelliStation
- MVS/ESA
- Netfinity
- NetVista
- Operating System/400
- OS/400
- RS/6000
- S/390
- Seascape
- SNAP/SHOT
- SP
- System/390
- System p5
- System Storage
- Versatile Storage Server
- Virtualization Engine
- VSE/ESA
- z/Architecture
- z/OS
- z/VM
- zSeries

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

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

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

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

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

Other company, product, and service names may be trademarks or service marks of others.

Terms and conditions

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

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

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

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

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

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

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

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 erfuehlt die Schutzanforderungen nach EN 50082-1 und EN 55022 Klasse A.

EN 55022 Klasse A Geraete beduerfen folgender Hinweise:

Nach dem EMVG:

"Geraete duerfen an Orten, fuer die sie nicht ausreichend entstoert sind, nur mit besonderer Genehmigung des Bundesministeriums fuer Post und Telekommunikation oder des Bundesamtes fuer Post und Telekommunikation

betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Störungen 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 Funkstörungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Massnahmen durchzuführen und dafür aufzukommen."

Anmerkung:

Um die Einhaltung des EMVG sicherzustellen, sind die Geräte wie in den Handbüchern 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

Index

Numerics

- 1750-511 and 1750-522 (DS6800) 4
- 1750-EX1 and 1750-EX2 expansion enclosures 5

A

- accessibility 118
- Activation (DSFA), Disk Storage Feature 14
- algorithms 85
- alias exposure 113
- API, DS Open 64
- array
 - arrays 11
 - definition 11
 - disk array 11

B

- balancing the I/O load 17
- base and alias exposures 113
- base exposure 113
- battery backup mode 97
- battery backup unit
 - indicators 56
 - overview 56

C

- Call Home support 91
- CCW, channel command words 111
- CKD, count key data storage 111
- CLI, command-line interface 65
- clusters, RAID disk groups 17
- concurrent code load 21
- configuration
 - DS series 71
 - DS Storage Manager 63
 - offline 73
 - online 73
 - real-time 73
 - reconfiguration 63
 - simulated 73
- Configuration overview 71
- connectivity, dial-up 91
- consistency groups
 - overview 31
- consolidating storage 107
- Copy Services
 - across 2105 and 1750 21
 - across 2105 and 2107 21
- count key data storage 111

D

- data availability and performance features 13
- data placement 17

- DDM
 - overview 43
- description 26
- determine storage configuration 7
- device driver, subsystem 15
- disaster recovery
 - fast restore 36
- disk array
 - definition 11
- disk drive module
 - indicators 43
 - overview 43
- disk drives
 - subsystem device driver 15
- Disk Manager
 - monitoring performance 65
- Disk Storage Feature Activation (DSFA) 14
- DS command-line interface 65
- DS Open API 64
- DS Storage Manager 63
- DS6000
 - major features 3
 - navigation 63
- DS6000 hardware resources 43
- DS6000 interfaces 63
- DSFA, Disk Storage Feature Activation 14
- dual power supply active mode 97

E

- enclosure identification 89
- expansion enclosure processor card 52
- express configuration
 - overview 73
- extent pool
 - overview 111
- extent pools
 - overview 111
- external communication 91, 103

F

- failback and failover 33
- fast reverse 36
- FATA disk drives 3
- fault analysis 99
- fault determination 99
- fault management 99
- FB, fixed block 112
- feature licensing 14
- features, data availability and performance 13
- fibre channel 39
- fibre-channel
 - basic topologies 40
 - host attachments 40
- fibre-channel adapters
 - attaching to open-systems host 41

- fibre-channel ATA (FATA) disk drives 3

FICON

- host systems 41
- multiple allegiance 113
- fixed block storage 112
- FlashCopy
 - changing direction of 25
 - incremental 24
 - multiple relationships 27
 - overview 24
 - persistent overview 27
 - refresh target volume 28
 - restrictions for open-systems hosts 39
- front display panel
 - indicators 45, 77
 - overview 45

G

- Global Copy
 - overview 30
- Global Mirror
 - description 34
 - key terms defined 34

H

- hardware resources 43
- host
 - attachments, fibre-channel 40
 - FICON attached zSeries 41
 - supported systems 39

I

- I/O load, balancing 17
- IBM service 103
- IBM System Storage Resiliency Family Copy Services 21
- introducing 21
- implementation, RAID 13
- incremental FlashCopy
 - refreshing
 - FlashCopy target 28
- indicators
 - battery backup unit 56
 - disk drive module 43
 - power supply 50
 - processor cards 52
 - rear operator panel 47
- interfaces, DS6000 63

K

- keyboards
 - accessibility features 118

L

- LCU
 - definition 112
- LED descriptions 77
- legal
 - terms and conditions 120
- licensed functions 14
- light path
 - overview 77
- light path overview 77
- location code 59
- locations, resources 59
- Logical control unit
 - definition 112
- logical subsystems
 - logical subsystems (LSSs)
 - overview 81
- logical volume manager
 - logical volume manager (LVM)
 - overview 83
- logical volumes 112
- LUN
 - calculation 85
- LVM (Logical Volume Manager) 83

M

- machine reported product data 91
- machine serial number 89
- machine type 89
- management console
 - multiple units 6
 - overview 6
- Metro Mirror
 - overview
 - write sequence 31
- Models 1750-511 and 1750-522 4
- Models 1750-EX1 and 1750-EX2 5
- MTMS label 89
- multiple allegiance facility 113
- multiple relationship FlashCopy 27

N

- near-line 3
- notification methods 91

O

- offline configuration overview 73
- online configuration overview 73
- open systems
 - overview 95
- overview
 - battery backup unit 56
 - configuring DS series 71
 - consistency groups 31
 - disk drive module 43
 - FlashCopy 24
 - Global Copy 30
 - offline configuration 73
 - online configuration 73
 - power supply 50
 - processor cards 52
 - real-time configuration 73

- overview (*continued*)
 - rear operator panel 47
 - Remote Mirror and Copy 29
 - sessions management 36
 - simulated configuration 73
 - system service card 58

P

- parallel access volume (PAV)
 - understanding static and dynamic 113
- path
 - connections 28
- PAV (parallel access volumes) 113
- performance features 9
- performance features, data
 - availability 13
- performance gathering
 - Disk Manager 65
- persistent FlashCopy
 - overview 27
- planning
 - licensed functions 14
- power control modes 97
- power supply
 - indicators 50
 - overview 50
- problem analysis 99
- problem determination 99
- problem management 99
- processor card
 - indicators 52
- processor cards
 - overview 52

R

- RAID
 - disk groups 17
 - implementation 13
 - RAID 10 overview 14
 - RAID 5 overview 13
 - RAID overview 13
- rank
 - definition 101
 - overview 101
- ranks
 - overview 101
- real-time configuration overview 73
- rear operator panel
 - indicators 47, 77
 - overview 47
- refresh volume, FlashCopy 24
- remote mirror and copy
 - restrictions for open-systems hosts 39
- Remote Mirror and Copy 29
- remote support 103
- Replication
 - copy services functions 66
- requirements
 - software 69
- resource location code 59
- resources
 - battery backup unit 56

- resources (*continued*)
 - disk drive module 43
 - fan 50
 - front display panel 45
 - power supply 50
 - processor cards 52
 - rear operator panel 47
- reversing
 - FlashCopy relationship 25

S

- SAN connections
 - fibre-channel adapters 41
- server enclosure processor card 52
- sessions, understanding 36
- simulated configuration overview 73
- single power supply active mode 97
- SNMP 105
- SNMP alerts 105
- SNMP support 91
- software requirements 69
- standby power mode 97
- storage, consolidating 107
- subsystem device driver (SSD) 15
- synchronous data mirroring 31
- system service card
 - overview 58

T

- TotalStorage Productivity Center
 - copy services 66
 - for Disk 65
 - Replication 66
- Trademarks 119

U

- understanding fixed block (FB)
 - architecture 112
- understanding logical volumes 112
- using FlashCopy
 - with remote mirror and copy 26

V

- virtual private network 115
- volume capacity
 - overview 85
- volume groups
 - overview 109

Z

- zones
 - host attachment 39
- zSeries hosts 41



Printed in USA