

IBM DS8870
Version 7 Release 5

Introduction and Planning Guide



Note

Before using this information and the product it supports, read the information in "Safety and environmental notices" on page vii and "Notices" on page 245.

This edition applies to version 7, release 5 of IBM DS8870 and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces GC27-4209-10.

© **Copyright IBM Corporation 2004, 2015.**

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

Contents

Safety and environmental notices vii

Safety notices and labels	vii
Caution notices for IBM DS8000 systems	viii
Danger notices for IBM DS8000 systems	xiii

About this book xvii

Who should use this book	xvii
Conventions and terminology	xvii
Publications and related information	xvii
Ordering IBM publications	xxi
Sending comments	xxi

Summary of changes xxiii

Chapter 1. Overview 1

Machine types overview	2
Hardware	3
Base frame (model 961) overview	4
Expansion frame (model 96E) overview	9
System types	11
Functional overview	14
Logical configuration	18
Logical configuration with DS8000 Storage Management GUI	18
Logical configuration with DS CLI	20
RAID implementation	22
Logical subsystems	24
Allocation methods	24
Management interfaces	25
DS8000 Storage Management GUI	26
DS command-line interface	26
DS Open Application Programming Interface	27
IBM Storage Mobile Dashboard	27
Tivoli Storage Productivity Center	28
Tivoli Storage Productivity Center for Replication	28
DS8000 Storage Management GUI supported web browsers	29

Chapter 2. Hardware features 31

Storage complexes	35
Management console	36
RAID implementation	36
RAID 5 overview	37
RAID 6 overview	37
RAID 10 overview	37
Hardware specifics	37
Storage system structure	37
Disk drives	38
Drive maintenance policy	38
Host attachment overview	39
Processor memory	41
Subsystem device driver for open-systems	41
I/O load balancing	42
Storage consolidation	42

Count key data	43
Fixed block	43
T10 DIF support	43
Logical volumes	44
Allocation, deletion, and modification of volumes	44
LUN calculation	45
Extended address volumes for CKD	46
Quick initialization	47

Chapter 3. Data management features 49

FlashCopy SE feature	49
Dynamic volume expansion	50
Count key data and fixed block volume deletion prevention	50
IBM Easy Tier	50
Easy Tier: automatic mode	53
Easy Tier: manual mode	61
Volume data monitoring	64
Easy Tier Heat Map Transfer Utility	64
Migration process management	67
Storage Tier Advisor tool	68
Easy Tier reporting improvements	68
Easy Tier considerations and limitations	68
VMware vStorage API for Array Integration support	69
Performance for IBM z Systems	71
Copy Services	72
Disaster recovery through Copy Services	81
Resource groups for Copy Services scope limiting	82
Comparison of licensed functions	84
I/O Priority Manager	85
Securing data	85

Chapter 4. Planning the physical configuration 87

Configuration controls	87
Determining physical configuration features	87
Management console features	88
Internal and external management consoles	88
Management console external power cord	89
Configuration rules for management consoles	90
Storage features	90
Storage enclosures and drives	90
Storage-enclosure fillers	95
Device adapters flash RAID adapters, and flash interface cards	96
Drive cables	96
Configuration rules for storage features	97
Physical and effective capacity	99
I/O adapter features	104
I/O enclosures	104
Fibre Channel (SCSI-FCP and FICON) host adapters and cables	105
Configuration rules for I/O adapter features	107
Processor complex features	111
Feature codes for processor licenses	111

Processor memory features	111
Feature codes for system memory	111
Configuration rules for system memory	112
Power features	112
Power cords	112
Input voltage	114
Direct-current uninterruptible-power supply	114
Configuration rules for power features	115
Other configuration features	115
Extended power line disturbance	115
Remote zSeries power control feature	116
BSMI certificate (Taiwan)	116
Shipping weight reduction	116

Chapter 5. Planning use of licensed functions 119

Licensed function indicators	119
License scope	120
Ordering licensed functions	122
Rules for ordering licensed functions	122
Operating environment license (239x Model LFA, OEL license, 242x machine type)	125
Feature codes for the operating-environment license	125
Parallel access volumes (239x Model LFA, PAV license; 242x machine type)	126
Feature codes for parallel access volume licensed function	127
IBM HyperPAV (242x Model PAV and 239x Model LFA, PAV license)	127
Feature code for IBM HyperPAV licensed function	127
IBM Easy Tier	128
Feature codes for IBM Easy Tier licensed function	128
Feature codes for IBM Easy Tier Server licensed function	128
Point-in-time copy function (239x Model LFA, PTC license) and FlashCopy SE Model SE function (239x Model LFA, SE license)	129
Feature codes for FlashCopy licensed function	129
Feature codes for Space Efficient FlashCopy licensed function	130
Remote mirror and copy functions (242x Model RMC and 239x Model LFA)	130
Feature codes for remote mirror and copy	130
Feature codes for I/O Priority Manager	132
z/OS licensed features	132
Remote mirror for z/OS (242x Model RMZ and 239x Model LFA, RMZ license)	132
Feature codes for z/OS Metro/Global Mirror	133
Incremental Resync licensed function	133
z/OS Distributed Data Backup	134
Thin provisioning licensed feature key	135

Chapter 6. Meeting delivery and installation requirements 139

Delivery requirements	139
Acclimation	139
Shipment weights and dimensions	139

Receiving delivery	140
Installation site requirements	141
Planning for floor and space requirements	141
Planning for power requirements	165
Planning for environmental requirements	173
Planning for safety	178
Planning for external management console installation	179
Planning for network and communications requirements	181

Chapter 7. Planning your storage complex setup 185

Company information	185
Management console network settings	185
Remote support settings	186
Notification settings	187
Power control settings	187
Control switch settings	187

Chapter 8. Planning data migration 191

Selecting a data migration method	192
---	-----

Chapter 9. Planning for security . . . 195

Planning for data encryption	195
Planning for encryption-key servers	195
Planning for key lifecycle managers	196
Planning for full-disk encryption activation	197
Planning for user accounts and passwords	197
Managing secure user accounts	197
Managing secure service accounts	197
Planning for NIST SP 800-131A security conformance	198

Chapter 10. License activation and management. 201

Planning your licensed functions	201
Activation of licensed functions	202
Activating licensed functions	202
Scenarios for managing licensing	204
Adding storage to your machine	204
Managing a licensed feature	204

Appendix A. Accessibility features for IBM DS8000 207

Appendix B. Warranty information 209

Appendix C. IBM DS8000 equipment and documents 211

Installation components	211
Customer components	212
Service components	212

Appendix D. DS8800 to DS8870 model conversion 213

DS8800 to DS8870 model conversion summary	213
Checking your preparations	215

Removing data, configuration, and encryption	216
Completing post-conversion tasks	216

**Appendix E. Customization
worksheets 217**

Company information worksheet	217
Management console network settings worksheet	219
Remote support worksheets	224
Outbound (call home and dump/trace offload) worksheet	224
Inbound (remote services) worksheets	229

Notification worksheets	234
SNMP trap notification worksheet	234
Email notification worksheet	235
Power control worksheet	237
Control switch settings worksheet	239

Notices 245

Trademarks	246
Homologation statement	247

Index 249

Safety and environmental notices

Review the safety notices, environmental notices, and electronic emission notices for this product before you install and use the product.

Safety notices and labels

Review the safety notices and safety information labels before using this product.

IBM Systems safety notices and information

This publication contains the safety notices for the IBM Systems products in English and other languages. It also contains the safety information labels found on the hardware in English and other languages. Anyone who plans, installs, operates, or services the system must be familiar with and understand the safety notices. Read the related safety notices before beginning work.



IBM Systems Safety Notices (www.ibm.com/shop/publications/order/), G229-9054

The publication is organized into three sections:

Safety notices

Lists the danger and caution notices without labels, organized alphabetically by language.

The following notices and statements are used in IBM documents. They are listed in order of decreasing severity of potential hazards.

Danger notice definition

A special note that calls attention to a situation that is potentially lethal or extremely hazardous to people.

Caution notice definition

A special note that calls attention to a situation that is potentially hazardous to people because of some existing condition, or to a potentially dangerous situation that might develop because of some unsafe practice.

Labels Lists the danger and caution notices that are accompanied with a label, organized by label reference number.

Text-based labels

Lists the safety information labels that might be attached to the hardware to warn of potential hazards, organized by label reference number.

Note: This product has been designed, tested, and manufactured to comply with IEC 60950-1, and where required, to relevant national standards that are based on IEC 60950-1.

Finding translated notices

Each safety notice contains an identification number. You can use this identification number to check the safety notice in each language. The list of notices that apply

to the product are listed in the “Danger notices for IBM DS8000 systems” on page xiii and the “Caution notices for IBM DS8000 systems” topics.

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

1. In the product documentation, look for the identification number at the end of each caution notice or each danger notice. In the following examples, the numbers (D002) and (C001) 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. (D002)

CAUTION:

A caution notice indicates the presence of a hazard that has the potential of causing moderate or minor personal injury. (C001)

2. Open the *IBM Systems Safety Notices* (G229-9054) publication.
3. Under the language, find the matching identification number. Review the topics concerning the safety notices to ensure that you are in compliance.

To view a PDF file, you need Adobe Reader. You can download it at no charge from the Adobe website (get.adobe.com/reader/).

Caution notices for IBM DS8000 systems

Ensure that you understand the caution notices for IBM® DS8000® systems.

Caution notices

Use the reference numbers in parentheses at the end of each notice, such as (C001), to find the matching translated notice in *IBM Systems Safety Notices*.

CAUTION: Energy hazard present. Shorting might result in system outage and possible physical injury. Remove all metallic jewelry before servicing. (C001)

CAUTION: Only trained service personnel may replace this battery. The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.
Do not: Throw or immerse into water, heat to more than 100°C (212°F), repair or disassemble. (C002)

CAUTION: Lead-acid batteries can present a risk of electrical burn from high, short circuit current. Avoid battery contact with metal materials; remove watches, rings, or other metal objects, and use tools with insulated handles. To avoid possible explosion, do not burn. (C004)

CAUTION: The battery is a lithium ion battery. To avoid possible explosion, do not burn. (C007)

CAUTION: The doors and covers to the product are to be closed at all times except for service by trained service personnel. All covers must be replaced and doors locked at the conclusion of the service operation. (C013)

CAUTION: The system contains circuit cards, assemblies, or both that contain lead solder. To avoid the release of lead (Pb) into the environment, do not burn. Discard the circuit card as instructed by local regulations. (C014)

CAUTION: This product is equipped with a 3-wire (two conductors and ground) power cable and plug. Use this power cable with a properly grounded electrical outlet to avoid electrical shock. (C018)

CAUTION: This product is equipped with a 4-wire (three-phase and ground) power cable. Use this power cable with a properly grounded electrical outlet to avoid electrical shock. (C019)

CAUTION: This product might be equipped with a 5-wire (three-phase, neutral ground) power cable. Use this power cable with a properly grounded electrical outlet to avoid electrical shock. (C020)

CAUTION: This product might be equipped with a hard-wired power cable. Ensure that a licensed electrician performs the installation per the national electrical code. (C022)

CAUTION: Ensure the building power circuit breakers are turned off **BEFORE** you connect the power cord or cords to the building power. (C023)

CAUTION: To avoid personal injury, disconnect the hot-swap, air-moving device cables before removing the fan from the device. (C024)

CAUTION: This assembly contains mechanical moving parts. Use care when servicing this assembly. (C025)

CAUTION: This product might contain one or more of the following devices: CD-ROM drive, DVD-ROM drive, DVD-RAM drive or laser module, which are Class 1 laser products. Note the following information:

- Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.
- Use of the controls or adjustments or performance of the procedures other than those specified herein might result in hazardous radiation exposure. (C026)

CAUTION: Servicing of this product or unit is to be performed by trained service personnel only. (C032)

CAUTION:



The weight of this part or unit is between 16 and 30 kg (35 and 66 lb). It takes two persons to safely lift this part or unit. (C040)

CAUTION: Refer to instruction manual. (C041)

CAUTION: Following the service procedure assures power is removed from 200-240VDC power distribution connectors before they are unplugged. However, unplugging 200-240VDC power distribution connectors while powered on, should not be done because it can cause connector damage and result in burn and /or shock injury from electrical arcing. (C043)



CAUTION: If your system has a module containing a lithium battery, replace it only with the same module type made by the same manufacturer. The battery contains lithium and can explode if not properly used, handled, or disposed of.

Do not:

- Throw or immerse into water
- Heat to more than 100°C (212°F)
- Repair or disassemble

Dispose of the battery as required by local ordinances or regulations. (C045)

Use the following general safety information for all rack mounted devices:

DANGER: Observe the following precautions when working on or around your IT rack system:

- Heavy equipment—personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices.



- Each rack cabinet might have more than one power cord. Be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (R001 part 1 of 2)

CAUTION:

- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.
- (For sliding drawers): Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.
- (For fixed drawers): This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack.
(R001 part 2 of 2)

CAUTION: Removing components from the upper positions in the rack cabinet improves rack stability during a relocation. Follow these general guidelines whenever you relocate a populated rack cabinet within a room or building.

- Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must observe the following precautions.
 - Remove all devices in the 32U position and above.
 - Ensure that the heaviest devices are installed in the bottom of the rack cabinet.
 - Ensure that there are no empty U-levels between devices installed in the rack cabinet below the 32U level.
- If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.
- Inspect the route that you plan to take to eliminate potential hazards.
- Verify that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.
- Verify that all door openings are at least 760 x 230 mm (30 x 80 in.).
- Ensure that all devices, shelves, drawers, doors, and cables are secure.
- Ensure that the four leveling pads are raised to their highest position.
- Ensure that there is no stabilizer bracket installed on the rack cabinet during movement.
- Do not use a ramp inclined at more than 10 degrees.
- When the rack cabinet is in the new location, complete the following steps:
 - Lower the four leveling pads.
 - Install stabilizer brackets on the rack cabinet.
 - If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.
- If a long-distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also lower the leveling pads to raise the casters off the pallet and bolt the rack cabinet to the pallet. (R002)

DANGER: Racks with a total weight of > 227 kg (500 lb.), Use Only Professional Movers! (R003)

DANGER: Do not transport the rack via fork truck unless it is properly packaged, secured on top of the supplied pallet. (R004)

CAUTION:

- Rack is not intended to serve as an enclosure and does not provide any degrees of protection required of enclosures.
- It is intended that equipment installed within this rack will have its own enclosure. (R005).

CAUTION: Use safe practices when lifting. (R007)


CAUTION: Do not place any object on top of a rack-mounted device unless that rack-mounted device is intended for use as a shelf. (R008)

DANGER:



Main Protective Earth (Ground):

This symbol is marked on the frame of the rack.

The PROTECTIVE EARTHING CONDUCTORS should be terminated at that point. A recognized or certified closed loop connector (ring terminal) should be used and secured to the frame with a lock washer using a bolt or stud. The connector should be properly sized to be suitable for the bolt or stud, the locking washer, the rating for the conducting wire used, and the considered rating of the breaker. The intent is to ensure the frame is electrically bonded to the PROTECTIVE EARTHING CONDUCTORS. The hole that the bolt or stud goes into where the terminal connector and the lock washer contact should be free of any non-conductive material to allow for metal to metal contact. All PROTECTIVE BONDING CONDUCTORS should terminate at this main protective earthing terminal or at points marked with  (R010)

Danger notices for IBM DS8000 systems

Ensure that you understand the danger notices for IBM DS8000 systems.

Danger notices

Use the reference numbers in parentheses at the end of each notice, such as (D001), to find the matching translated notice in *IBM Systems Safety Notices*.

DANGER: To prevent a possible shock from touching two surfaces with different protective ground (earth), use one hand, when possible, to connect or disconnect signal cables. (D001)

DANGER: Overloading a branch circuit is potentially a fire hazard and a shock hazard under certain conditions. To avoid these hazards, ensure that your system electrical requirements do not exceed branch circuit protection requirements. Refer to the information that is provided with your device or the power rating label for electrical specifications. (D002)

DANGER: An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (D004)

DANGER: When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- Connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
 2. Attach all cables to the devices.
 3. Attach the signal cables to the connectors.
 4. Attach the power cords to the outlets.
 5. Turn on the devices.
- Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

DANGER: Heavy equipment — personal injury or equipment damage might result if mishandled. (D006)

DANGER: Uninterruptible power supply (UPS) units contain specific hazardous materials. Observe the following precautions if your product contains a UPS:

- The UPS contains lethal voltages. All repairs and service must be performed only by an authorized service support representative. There are no user serviceable parts inside the UPS.
- The UPS contains its own energy source (batteries). The output receptacles might carry live voltage even when the UPS is not connected to an AC supply.
- Do not remove or unplug the input cord when the UPS is turned on. This removes the safety ground from the UPS and the equipment connected to the UPS.
- The UPS is heavy because of the electronics and batteries that are required. To avoid injury, observe the following precautions:
 - Do not attempt to lift the UPS by yourself. Ask another service representative for assistance.
 - Remove the battery, electronics assembly, or both from the UPS before removing the UPS from the shipping carton or installing or removing the UPS in the rack. (D007)

DANGER: Professional movers are to be used for all relocation activities. Serious injury or death may occur if systems are handled and moved incorrectly. (D008)

About this book

This book describes how to plan for a new installation of DS8870. It includes information about planning requirements and considerations, customization guidance, and configuration worksheets.

Who should use this book

This book is intended for personnel that are involved in planning. Such personnel include IT facilities managers, individuals responsible for power, cooling, wiring, network, and general site environmental planning and setup.

Conventions and terminology

Different typefaces are used in this guide to show emphasis, and various notices are used to highlight key information.

The following typefaces are used to show emphasis:

Typeface	Description
Bold	Text in bold represents menu items.
bold monospace	Text in bold monospace represents command names.
<i>Italics</i>	Text in italics is used to emphasize a word. In command syntax, it is used for variables for which you supply actual values, such as a default directory or the name of a system.
Monospace	Text in monospace identifies the data or commands that you type, samples of command output, examples of program code or messages from the system, or names of command flags, parameters, arguments, and name-value pairs.

These notices are used to highlight key information:

Notice	Description
Note	These notices provide important tips, guidance, or advice.
Important	These notices provide information or advice that might help you avoid inconvenient or difficult situations.
Attention	These notices indicate possible damage to programs, devices, or data. An attention notice is placed before the instruction or situation in which damage can occur.

Publications and related information

Product guides, other IBM publications, and websites contain information that relates to the IBM DS8000 series.

To view a PDF file, you need Adobe Reader. You can download it at no charge from the Adobe website (get.adobe.com/reader/).

Online documentation

The IBM DS8000 series online product documentation (www.ibm.com/support/knowledgecenter/ST8NCA/product_welcome/ds8000_kcwelcome.html) contains all of the information that is required to install, configure, and manage DS8000 storage systems. The online documentation is updated between product releases to provide the most current documentation.

Publications

You can order or download individual publications (including previous versions) that have an order number from the IBM Publications Center website (www.ibm.com/shop/publications/order/). Publications without an order number are available on the documentation CD or can be downloaded here.

Table 1. DS8000 series product publications

Title	Description	Order number
<i>DS8870 Introduction and Planning Guide</i>	This publication provides an overview of the product and technical concepts for DS8870. It also describes the ordering features and how to plan for an installation and initial configuration of the storage system.	V7.5 GC27-4209-11 V7.4 GC27-4209-10 V7.3 GC27-4209-09 V7.2 GC27-4209-08 V7.1 GC27-4209-05 V7.0 GC27-4209-02
<i>DS8800 and DS8700 Introduction and Planning Guide</i>	This publication provides an overview of the product and technical concepts for DS8800 and DS8700. It also describes ordering features and how to plan for an installation and initial configuration of the storage system.	V6.3 GC27-2297-09 V6.2 GC27-2297-07
<i>Host Systems Attachment Guide</i>	This publication provides information about attaching hosts to the storage system. You can use various host attachments to consolidate storage capacity and workloads for open systems and IBM z Systems™ hosts.	V7.5 GC27-4210-04 V7.4 GC27-4210-03 V7.2 GC27-4210-02 V7.1 GC27-4210-01 V7.0 GC27-4210-00 V6.3 GC27-2298-02
<i>IBM Storage System Multipath Subsystem Device Driver User's Guide</i>	This publication provides information regarding the installation and use of the Subsystem Device Driver (SDD), Subsystem Device Driver Path Control Module (SDDPCM), and Subsystem Device Driver Device Specific Module (SDDDSM) on open systems hosts.	Download
<i>Command-Line Interface User's Guide</i>	This publication describes how to use the DS8000 command-line interface (DS CLI) to manage DS8000 configuration and Copy Services relationships, and write customized scripts for a host system. It also includes a complete list of CLI commands with descriptions and example usage.	V7.5 GC27-4212-06 V7.4 GC27-4212-04 V7.3 GC27-4212-03 V7.2 GC27-4212-02 V7.1 GC27-4212-01 V7.0 GC27-4212-00 V6.3 GC53-1127-07
<i>Application Programming Interface Reference</i>	This publication provides reference information for the DS8000 Open application programming interface (DS Open API) and instructions for installing the Common Information Model Agent, which implements the API.	V7.3 GC27-4211-03 V7.2 GC27-4211-02 V7.1 GC27-4211-01 V7.0 GC35-0516-10 V6.3 GC35-0516-10

Table 1. DS8000 series product publications (continued)

Title	Description	Order number
<i>RESTful API Guide</i>	This publication provides an overview of the Representational State Transfer (RESTful) API, which provides a platform independent means by which to initiate create, read, update, and delete operations in the DS8000 and supporting storage devices.	V1.0 SC27-8502-00

Table 2. DS8000 series warranty, notices, and licensing publications

Title	Order number
<i>Warranty Information for DS8000 series</i>	See the DS8000 Publications CD
<i>IBM Safety Notices</i>	Search for G229-9054 on the IBM Publications Center website
<i>IBM Systems Environmental Notices</i>	http://ibm.co/1fBgWFI
<i>International Agreement for Acquisition of Software Maintenance (Not all software will offer Software Maintenance under this agreement.)</i>	http://ibm.co/1fBmKPz
<i>License Agreement for Machine Code</i>	http://ibm.co/1mNiW1U
<i>Other Internal Licensed Code</i>	http://ibm.co/1kvABXE
<i>International Program License Agreement and International License Agreement for Non-Warranted Programs</i>	www.ibm.com/software/sla/sladb.nsf/pdf/ilan/\$file/ilan_en.pdf

See the Agreements and License Information CD that was included with the DS8000 series for the following documents:

- License Information
- Notices and Information
- Supplemental Notices and Information

Related publications

Listed in the following table are the IBM Redbooks® publications, technical papers, and other publications that relate to DS8000 series.

Table 3. DS8000 series related publications

Title	Description
IBM Security Key Lifecycle Manager online product documentation (www.ibm.com/support/knowledgecenter/SSWPVP/)	This online documentation provides information about IBM Security Key Lifecycle Manager, which you can use to manage encryption keys and certificates.

Table 3. DS8000 series related publications (continued)

Title	Description
IBM Tivoli® Storage Productivity Center online product documentation (www.ibm.com/support/knowledgecenter/SSNE44/)	This online documentation provides information about Tivoli Storage Productivity Center, which you can use to centralize, automate, and simplify the management of complex and heterogeneous storage environments including DS8000 storage systems and other components of your data storage infrastructure.

Related websites

View the websites in the following table to get more information about DS8000 series.

Table 4. DS8000 series related websites

Title	Description
IBM website (ibm.com®)	Find more information about IBM products and services.
IBM Support Portal website (www.ibm.com/storage/support)	Find support-related information such as downloads, documentations, troubleshooting, and service requests and PMRs.
IBM Directory of Worldwide Contacts website (www.ibm.com/planetwide)	Find contact information for general inquiries, technical support, and hardware and software support by country.
IBM DS8000 series website (www.ibm.com/servers/storage/disk/ds8000)	Find product overviews, details, resources, and reviews for the DS8000 series.
IBM System Storage® Interoperation Center (SSIC) website (www.ibm.com/systems/support/storage/config/ssic)	Find information about host system models, operating systems, adapters, and switches that are supported by the DS8000 series.
IBM Storage SAN (www.ibm.com/systems/storage/san)	Find information about IBM SAN products and solutions, including SAN Fibre Channel switches.
IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa)	Download licensed machine code (LMC) feature keys that you ordered for your DS8000 storage systems.
IBM Fix Central (www-933.ibm.com/support/fixcentral)	Download utilities such as the IBM Easy Tier® Heat Map Transfer utility and Storage Tier Advisor tool.
IBM Java™ SE (JRE) (www.ibm.com/developerworks/java/jdk)	Download IBM versions of the Java SE Runtime Environment (JRE), which is often required for IBM products.
DS8700 Code Bundle Information website (www.ibm.com/support/docview.wss?uid=ssg1S1003593)	Find information about code bundles for DS8700. See section 3 for web links to SDD information. The version of the currently active installed code bundle now displays with the DS CLI ver command when you specify the -l parameter.

Table 4. DS8000 series related websites (continued)

Title	Description
DS8800 Code Bundle Information website(www.ibm.com/support/docview.wss?uid=ssg1S1003740)	Find information about code bundles for DS8800. See section 3 for web links to SDD information. The version of the currently active installed code bundle now displays with the DS CLI <code>ver</code> command when you specify the <code>-l</code> parameter.
DS8870 Code Bundle Information website (www.ibm.com/support/docview.wss?uid=ssg1S1004204)	Find information about code bundles for DS8870. See section 3 for web links to SDD information. The version of the currently active installed code bundle now displays with the DS CLI <code>ver</code> command when you specify the <code>-l</code> parameter.

Ordering IBM publications

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

Procedure

The IBM Publications Center website (www.ibm.com/shop/publications/order/) offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download at no charge. You can also order publications. The IBM Publications Center website displays prices in your local currency.

Sending comments

Your feedback is important in helping to provide the most accurate and highest quality information.

Procedure

To submit any comments about this publication or any other IBM storage product documentation:

Send your comments by email to starpubs@us.ibm.com. Be sure to include the following information:

- Exact publication title and version
- Publication form number (for example, GA32-1234-00)
- Page, table, or illustration numbers that you are commenting on
- A detailed description of any information that should be changed

Summary of changes

IBM DS8870 Version 7, Release 5 introduces the following new features.

Version 7.5

This table provides the current technical changes and enhancements to the IBM DS8870. Changed and new information is indicated by a vertical bar (|) to the left of the change.

Function	Description
Support for model AP1 External Security Key Life-cycle Manager (SKLM) isolated-key appliance single or dual processor configurations	See “Machine types overview” on page 2, Chapter 2, “Hardware features,” on page 31, “DS8000 support appliances” on page 172, and “Planning for key lifecycle managers” on page 196 for more information.
Support for 4-port, 16 Gbps shortwave and longwave FCP and FICON® host adapter, PCIe	See “Feature codes for Fibre Channel host adapters” on page 105 for more information.
T10 DIF volume active protection support for AIX® on IBM Power Systems™	See “T10 DIF support” on page 43 for more information.
Support for IBM z Systems synergy	DS8870 and IBM z Systems synergy improves performance, availability, and growth management. The control switch settings worksheet is updated with four new control switches. See “Control switch settings worksheet” on page 239 for more information.
Heat Map Transfer for Metro Global Mirror (MGM) environments	The heat map transfer utility supports MGM replication in addition to Metro Mirror, Global Copy, and Global Mirror functions. See “Easy Tier Heat Map Transfer Utility” on page 64 for more information.
Remote user directory policy	You can use a local administrator account to access a DS8000 system when a remote user directory policy is configured and the remote user directory server is inaccessible. See the setauthpol command that is provided in Knowledge Center. To view the information in Knowledge Center, use the search or filtering functions, or find it in the navigation by clicking System Storage > Disk systems> Enterprise Storage Servers> DS8000 . Go to IBM Knowledge Center website to learn more.

Chapter 1. Overview

IBM DS8870 is a high-performance, high-capacity storage system that supports continuous operation, data security, and data resiliency. It is the latest and most advanced storage system in the IBM DS8000 series. The storage system consists of a base frame (model 961), optionally up to three expansion frames (model 96E), and one or two management consoles (two being the recommended configuration). For high-availability, the hardware components are redundant.

DS8870 adds base frame and expansion frame to the 242x machine type family.

- The base frame contains the processor nodes, I/O enclosures, Ethernet switches, and the Hardware Management Console (HMC), in addition to power and storage enclosures. The base frame is available with different processor options that range from dual two-core systems to dual 16-core systems.
- Depending on the system configuration, you can add up to three expansion frames to the storage system. Only the first expansion frame contains I/O enclosures, which provide more host adapters, device adapters, and flash RAID adapters.

An optional external HMC is recommended for high-availability.

DS8870 integrates high-performance flash enclosures and flash cards to provide a higher level of performance for DS8870. The flash enclosures and flash cards are supported in the Enterprise Class, Business Class, and all Flash configurations. The DS8870 All-Flash configuration provides twice the I/O bays and up to twice the host adapters as the standard DS8870 single-frame configuration. DS8870 continues to be available in a standard configuration with disk drives and flash drives in the Enterprise Class and Business Class configurations.

DS8870 also includes the following features such as:

- POWER7+™ processors
- Power-usage reporting
- National Institute of Standards and Technology (NIST) SP 800-131A enablement

Other functions that are supported in both the DS8000 Storage Management GUI and the DS command-line interface (DS CLI) include:

- Easy Tier
- Data encryption
- Thin provisioning

You can use the DS8000 Storage Management GUI and the DS command-line interface (DS CLI) to manage and logically configure the storage system.

Functions that are supported in only the DS command-line interface (DS CLI) include:

- Point-in-time copy functions with IBM FlashCopy®
- Space Efficient FlashCopy
- Remote Mirror and Copy functions, including
 - Metro Mirror
 - Global Copy
 - Global Mirror
 - Metro/Global Mirror

- z/OS® Global Mirror
- z/OS Metro/Global Mirror
- Multiple Target PPRC
- I/O Priority Manager

DS8870 meets hazardous substances (RoHS) requirements by conforming to the following EC directives:

- Directive 2011/65/EU of the European Parliament and of the Council of 8 June 2011 on the restriction of the use of certain hazardous substances in electrical and electronic equipment. It has been demonstrated that the requirements specified in Article 4 are met.
- EN 50581:2012 technical documentation for the assessment of electrical and electronic products regarding the restriction of hazardous substances.

The IBM Security Key Lifecycle Manager (formerly known as Tivoli Key Lifecycle Manager) stores data keys that are used to secure the key hierarchy that is associated with the data encryption functions of various devices, including the DS8000 series. It can be used to provide, protect, and maintain encryption keys that are used to encrypt information that is written to and decrypt information that is read from encryption-enabled disks. IBM Security Key Lifecycle Manager operates on various operating systems.

Note: You can convert a DS8800 series to a DS8870 enterprise-class configuration.

Machine types overview

Several machine type options are available. Order a hardware machine type for the storage system and a corresponding function authorization machine type for the licensed functions that are planned for use.

The following table lists the available hardware machine types and their corresponding function authorization machine types.

Table 5. Available hardware and function-authorization machine types

Hardware		Licensed functions	
Hardware machine type	Available hardware models	Corresponding function authorization machine type	Available function authorization models
2421 (1-year warranty period)	961 or AP1, and 96E	2396 (1-year warranty period)	LFA
2422 (2-year warranty period)	961, 96E	2397 (2-year warranty period)	
2423 (3-year warranty period)		2398 (3-year warranty period)	
2424 (4-year warranty period)		2399 (4-year warranty period)	

An intermix of 242x hardware machine types (warranty machine types) is supported within one storage system. For example, you can have a storage system that is composed of a 2421 model 961 (one-year warranty) and a 2423 model 96E (three-year warranty).

Because the 242x hardware machine types are built on the 2107 machine type and microcode, some interfaces might display 2107. This display is normal and is no cause for alarm. The 242x machine type that you purchased is the valid machine type.

Hardware

The architecture of the IBM DS8000 is based on three major elements that provide function specialization and three tiers of processing power.

Figure 1 illustrates the following elements.

- Host adapters manage external I/O interfaces that use Fibre Channel protocols for host-system attachment and for replicating data between storage systems.
- Flash RAID adapters and device adapters manage the internal storage devices. They also manage the SAS paths to drives, RAID protection, and drive sparing.
- A pair of high-performance redundant active-active Power® servers is functionally positioned between the adapters and a key feature of the architecture.

The internal Power servers support the bulk of the processing to be done in the storage system. Each Power server has multiple processor cores. The cores are managed as a symmetric multiprocessing (SMP) pool of shared processing power to process the work that is done on the Power server. Each Power server runs an AIX kernel that manages the processors, manages processor memory as a data cache, and more. For more information, see *IBM DS8000 Architecture and Implementation* on the IBM Redbooks website (www.redbooks.ibm.com/).

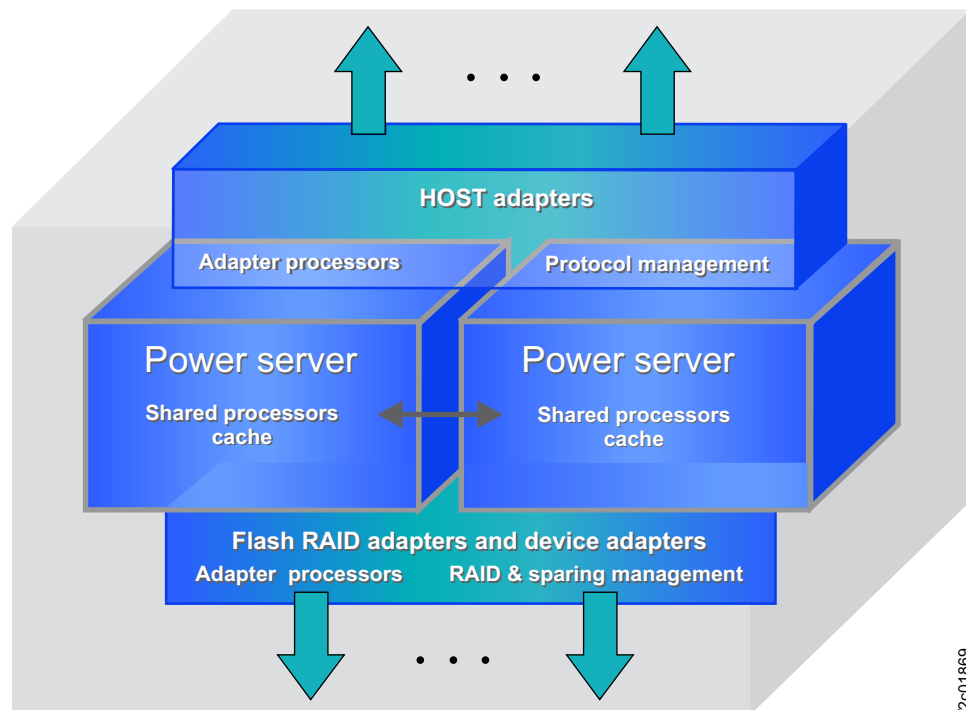


Figure 1. DS8000 architecture

The DS8000 architecture has the following major benefits.

- Server foundation

- Promotes high availability and high performance by using field-proven Power servers
- Reduces custom components and design complexity
- Positions the storage system to reap the benefits of server technology advances
- Operating environment
 - Promotes high availability and provides a high-quality base for the storage system software through a field-proven AIX operating-system kernel
 - Provides an operating environment that is optimized for Power servers, including performance and reliability, availability, and serviceability
 - Provides shared processor (SMP) efficiency
 - Reduces custom code and design complexity
 - Uses Power firmware and software support for networking and service functions

Base frame (model 961) overview

DS8870 includes a base frame (model 961).

The base frame includes the following components:

- High-performance flash enclosures **7**
- Standard drive enclosures **1**
- Ethernet switches **2**
- Internal Hardware Management Console (HMC) **3**
- External HMC (optional)
- Processor nodes (available with POWER7+ and POWER7[®] processors) **4**
- I/O enclosures **5**
- Direct-current uninterruptible power supplies (DC-UPS) **6**
- Rack power control (RPC) cards **8** (The RPC cards are visible from the back of the base frame.)

Figure 2 on page 5 illustrates an example of a *standard configuration* for a base frame with the maximum number of standard disk enclosures and flash enclosures. Figure 3 on page 6 illustrates an example of an *all-flash configuration* for a base frame with the maximum number flash enclosures.




Figure 2. DS8870 base frame with a standard configuration

2c02129




Figure 3. DS8870 base frame with an all-flash configuration

Storage enclosures

DS8870 integrate one of two types of storage enclosures: high-performance flash enclosures and standard drive enclosures.

High-performance flash enclosures:

The high-performance flash enclosure is a 1U RAID storage enclosure that is installed individually. It is not installed in pairs, as are standard drive enclosures.

A DS8870 with standard drive enclosures supports up to four high-performance flash enclosures in a base frame (model 961) and up to four high-performance flash enclosures in a first expansion frame (model 96E).

A DS8870 All-Flash configuration can contain up to eight high-performance flash enclosures (four vertically and four horizontally) in a single base frame.

Note: DS8870 storage systems that are converted from DS8800 storage systems are either all non-FDE or all FDE at the time of conversion. However, with release 7.4 or later, high-performance flash enclosures can be added to a converted DS8870 non-FDE system.

Each high-performance flash enclosure contains the following hardware components:

- 16 or 30 400-GB 1.8-inch SAS flash cards, which support IBM Full Disk Encryption (FDE)
- Two power supplies with integrated cooling fans
- Two Flash RAID adapters, which are configured as a pair, that provide redundant data path to the flash cards in the high-performance flash enclosure. These adapters also provide enclosure control.
- One back plane for plugging components

The high-performance flash enclosures connect to the I/O enclosures over a PCI Express® (PCIe) fabric, which increases bandwidth and transaction-processing capability. The I/O enclosures connect to POWER7+ processor complexes over a PCIe bus fabric.

The RAID controller is designed to unleash the performance capabilities of flash-based storage. Currently, the flash cards support RAID 5 arrays and mirrored/protected write cache.

Restriction:

- The flash cards are not available as Standby CoD drives.
- The flash cards are not supported in RAID-6 or RAID-10 configurations.




Figure 4. Flash enclosure (front and rear views)

Standard drive enclosures:

The standard drive enclosure is a 2U storage enclosure that is installed in pairs.

Each standard drive enclosure contains the following hardware components:

- Up to 12 large-form factor (LFF), 3.5-inch drive enclosures
- Up to 24 small form factor (SFF), 2.5-inch SAS drives

Note: Drives can be disk drives or flash drives (also known as solid-state drives or SSDs). You cannot intermix drives of different types in the same enclosure.

- Two power supplies with integrated cooling fans
- Two Fibre Channel interconnect cards that connect four Fibre Channel 8 Gbps interfaces to a pair of device adapters or another standard drive enclosure.

- One back plane for plugging components

The 2.5-inch disk drives are available in sets of 16 drives. The 3.5-inch SAS disk drives are available in half-drive sets of eight drives. Flash drives are available in sets of 16 or half-drive sets of eight drives.

Disk drives are available as Standby capacity on demand (Standby CoD). Using the Standby CoD features, you can install inactive drives that can be easily activated as business needs require. The storage system offers up to six Standby CoD disk drive sets that can be factory-installed or field-installed. To activate the Standby CoD disk drive set, you logically configure the disk drives for use. Activation is a nondisruptive activity and does not require intervention from IBM. After any portion of the Standby CoD disk drive set is activated, you must place an order with IBM to initiate billing for the activated set. You can also order replacement Standby CoD disk drive sets.

Note: The flash drives are not available as Standby CoD drives.

Management console

The management console (or *management server*) is also referred to as the Hardware Management Console (or HMC). It supports storage system hardware and firmware installation and maintenance activities. The HMC includes a keyboard and display.

The HMC connects to the customer network and provides access to IBM DS8000 functions that can be used to manage the storage system. Management functions include logical configuration, problem notification, call home for service, remote service, and Copy Services management. You can perform management functions from the DS8000 Storage Management GUI, DS command-line interface (DS CLI), or other storage management software that supports the DS8000 series.

One HMC is physically located inside the storage system. A second external HMC is available as a separately orderable feature to provide redundancy.

Ethernet switches

The Ethernet switches provide internal communication between the management consoles and the processor complexes. Two redundant Ethernet switches are provided.

Processor nodes

The processor nodes drive all functions in the storage system. Each node consists of a Power server that contain POWER7 or POWER7+ processors and memory.

The POWER7+ processor delivers performance improvements in I/O operations in transaction processing workload over the previous POWER7 processor.

I/O enclosures

I/O enclosures provide connectivity between the adapters and the processor complex.

The I/O enclosure uses PCIe interfaces to interconnect I/O adapters in the I/O enclosure to both processor nodes. A PCIe device is an I/O adapter or a processor node.

To improve I/O operations per second (IOPS) and sequential read/write throughput, each I/O enclosure is connected to each processor node with a point-to-point connection. I/O enclosures no longer share common loops.

I/O enclosures contain the following adapters:

Flash interface cards

Interface card that provides PCIe cable connection from the I/O enclosure to the high-performance flash enclosure.

Device adapters

PCIe-attached adapter with four 8 Gbps Fibre Channel arbitrated loop (FC-AL) ports. These adapters connect the processor nodes to standard drive enclosures and provide RAID support.

Host adapters

PCIe-attached adapter with four or eight 8 Gbps Fibre Channel ports. Both longwave and shortwave adapter versions that support different maximum cable lengths are available. Each port can be independently configured to use SCSI/FCP, SCSI/FC-AL, or FICON/FCX protocols. The host-adapter ports can be directly connected to attached hosts systems or storage systems, or connected to a storage area network. SCSI/FCP ports are used for connections between storage systems. SCSI/FCP ports that are attached to a SAN can be used for both host and storage system connections.

The High Performance FICON Extension (FCX) protocol can be used by FICON host channels that have FCX support. The use of FCX protocols provides a significant reduction in channel usage. This reduction improves I/O input on a single channel and reduces the number of FICON channels that are required to support the workload.

Power

The power system in each frame is a pair of direct-current uninterruptible power supplies (DC-UPSs) with internal batteries. The DC-UPSs distribute rectified ac power and provide power switching for redundancy. A single DC-UPS has sufficient capacity to power and provide battery backup to the entire frame if one DC-UPS is out of service.

There are two ac-power cords, each feeding one DC-UPS. If ac power is not present at the input line, the output is switched to rectified ac power from the partner DC-UPS. If neither ac-power input is active, the DC-UPS switches to 208 V dc battery power. Storage systems that have the extended power line disturbance (ePLD) option are protected from a power-line disturbance for up to 50 seconds. Storage systems without the ePLD option are protected for 4 seconds.

An integrated pair of rack-power control (RPC) cards manages the efficiency of power distribution within the storage system. The RPC cards are attached to each processor node. The RPC card is also attached to the primary power system in each frame.

Expansion frame (model 96E) overview

An expansion frame (model 96E) is supported for DS8870 Enterprise Class and Business Class configurations with a minimum of 128 GB system memory. The DS8870 All-Flash configuration does not support an expansion frame.

For DS8870 Enterprise Class, up to three expansion frames can be added to a base frame with a supported configuration. The first expansion frame supports up to 336 2.5-inch disk drives. The second expansion frame supports up to 480 2.5-inch

disk drives. A third expansion frame supports an extra 480 2.5-inch disk drives. When all four frames are used, DS8870 can support a total of 1,536 2.5-inch disk drives in a compact footprint, creating a high-density storage system, preserving valuable floor space in data center environments, and reducing power consumption.

Only the first expansion frame includes I/O enclosures. You can add up to four additional high-performance flash enclosures to the first expansion frame. See Figure 5. The second and third expansion frames do not include I/O enclosures. See Figure 6 on page 11.

The main power area is at the rear of the expansion frame. The power system in each frame is a pair of direct-current uninterruptible power supplies (DC-UPSs) with internal batteries.

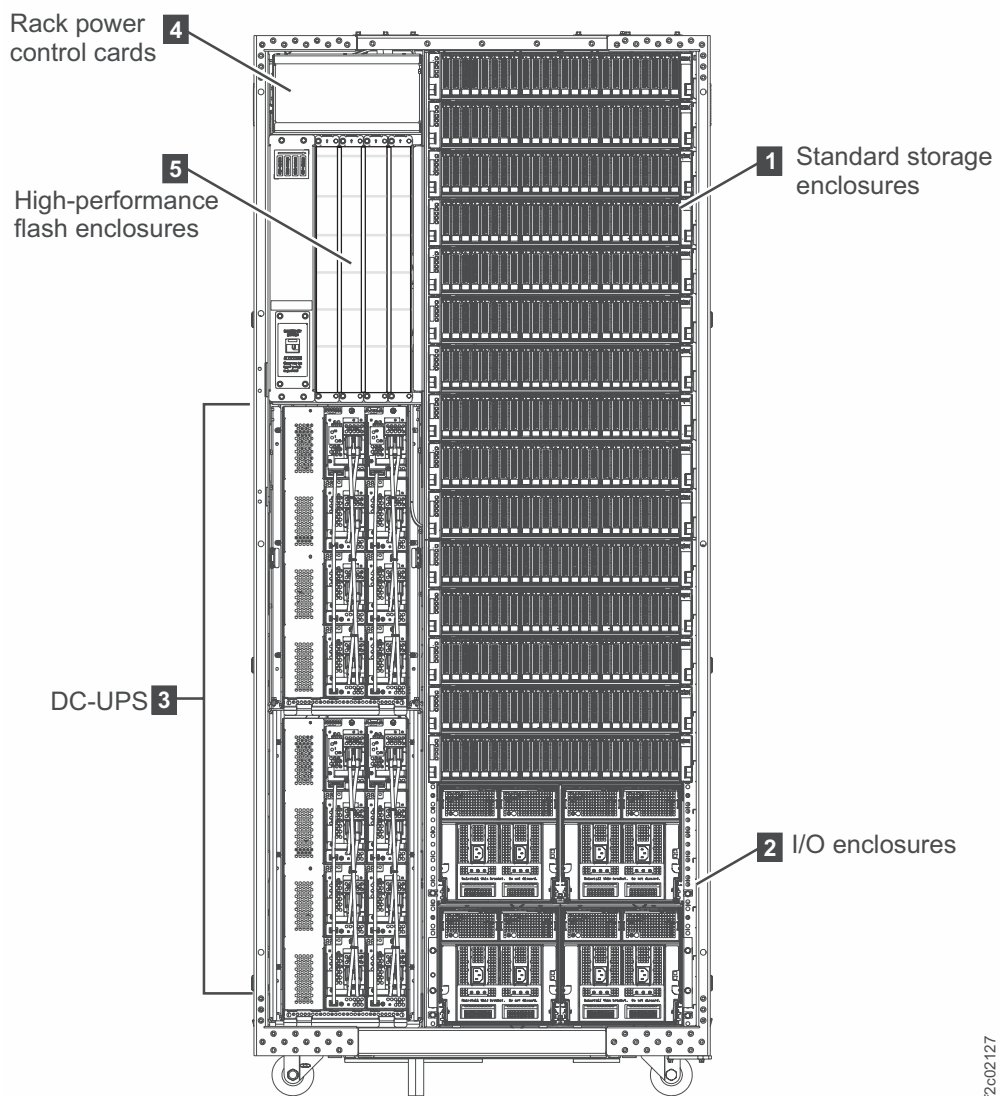


Figure 5. Four high-performance flash enclosures in an expansion frame

12c02127




Figure 6. DS8870 second and third expansion frame

12c02081

System types

DS8870 supports three configurations: All Flash, Enterprise Class, and Business Class.

For more specifications, see the IBM DS8000 series specifications web site (www.ibm.com/systems/storage/disk/ds8000/specifications.html).

All Flash configurations

The DS8870 All Flash configuration is a high-performance configuration that supports up to eight high-performance flash enclosures. Flash enclosures 1 - 4 are mounted vertically, and flash enclosures 5 - 8 are mounted horizontally.

This configuration does not support expansion frames (model 96E). This configuration also does not support standard drive enclosures with disk drives or flash drives (SSDs).

This configuration uses 8 or 16 Gbps Fibre Channel host adapters that run Fibre Channel Protocol (FCP), FICON, or Fibre Channel Arbitrated Loop (FC-AL) protocol. The High Performance FICON (HPF) feature is also supported.

The following tables list the hardware components and maximum capacities that are supported for the All Flash configuration, depending on the amount of memory that is available.

Table 6. Components and maximum capacity for All Flash configurations

Processor	Total system memory	Processor memory	I/O enclosures	Flash RAID adapter pairs	Host adapters (8 or 4 port)	Flash enclosures	Maximum flash cards	Maximum storage capacity for 1.8-in. flash cards	Expansion frames
8-core	256 GB	128 GB	8	1 - 8	2 - 16	1 - 8	240	96 TB	0
16-core	512 GB	256 GB	8	1 - 8	2 - 16	1 - 8	240	96 TB	0
16-core	1024 GB	512 GB	8	1 - 8	2 - 16	1 - 8	240	96 TB	0

Enterprise Class configurations

The Enterprise Class configuration is a high-density, high-performance configuration that includes standard disk enclosures and high-performance flash enclosures.

Enterprise Class storage systems are scalable up to 16-core processors, with up to 240 flash cards, and up to 1,536 standard drives. They are optimized and configured for performance and throughput, by maximizing the number of device adapters and paths to the storage enclosures. They support the following storage enclosures:

- Up to 5 standard drive enclosure pairs and up to four high-performance flash enclosures in a base frame (model 961).
- Up to 7 standard drive enclosures pairs and up to 4 high-performance flash enclosures in a first expansion frame (model 96E).
- Up to 10 standard drive enclosures pairs in a second expansion frame.
- Up to 10 standard drive enclosures pairs in a third expansion frame.

Enterprise Class uses 8 or 16 Gbps Fibre Channel host adapters that run Fibre Channel Protocol (FCP), FICON, or Fibre Channel Arbitrated Loop (FC-AL) protocol. The High Performance FICON (HPF) feature is also supported.

This configuration supports three-phase and single-phase power.

Restriction: Copy Services and I/O Priority Manager functions require a minimum of 32 GB system memory.

For more specifications, see the IBM DS8000 series specifications web site (www.ibm.com/systems/storage/disk/ds8000/specifications.html).

The following tables list the hardware components and maximum capacities that are supported for the business-class configuration, depending on the amount of memory that is available.

Table 7. Components for the Enterprise Class configuration

Processors	System memory	Processor memory	I/O enclosures	Host adapters (8 or 4 port)	Device adapter pairs	Flash RAID adapter pairs	Standard drive enclosure pairs ^{1, 2}	Flash enclosures ²	Expansion frames
2-core	16 GB	8 GB	2	2 - 4	1 - 2	0	1 - 3	0	0
2-core	32 GB	16 GB	2	2 - 4	1 - 2	0 - 2	0 - 3	0 - 2	0
4-core	64 GB	32 GB	4	2 - 8	1 - 4	0 - 4	0 - 5	0 - 4	0
8-core	128 GB	64 GB	8	2 - 16	1 - 8	0 - 8	0 - 22	0 - 8	0 - 2
8-core	256 GB	128 GB	8	2 - 16	1 - 8	0 - 8	0 - 32	0 - 8	0 - 3
16-core	512 GB	256 GB	8	2 - 16	1 - 8	0 - 8	0 - 32	0 - 8	0 - 3
16-core	1,024 GB	512 GB	8	2 - 16	1 - 8	0 - 8	0 - 32	0 - 8	0 - 3

1. Standard drive enclosures are installed in pairs.
 2. This configuration of the DS8870 must be populated with either one standard drive enclosure pair (feature code 1241) or one high-performance flash enclosure (feature code 1500).

Table 8. Maximum capacity for the Enterprise Class configuration

Processors	System memory	Maximum 2.5-in. disk drives	Maximum storage capacity for 2.5-in. disk drives	Maximum 3.5-in. disk drives	Maximum storage capacity for 3.5-in. disk drives	Maximum 1.8-in. flash cards	Maximum storage capacity for 1.8-in. flash cards	Maximum total drives ¹
2-core	16 GB	144	230.4 TB	72	288 TB	N/A	N/A	144
2-core	32 GB	144	230.4 TB	72	288 TB	60	24 TB	204
4-core	64 GB	240	384 TB	120	480 TB	120	48 TB	360
8-core	128 GB	1056	1.69 PB	528	2.1 PB	240	96 TB	1296
8-core	256 GB	1536	2.46 PB	768	3 PB	240	96 TB	1776
16-core	512 GB	1536	2.46 PB	768	3 PB	240	96 TB	1776
16-core	1,024 GB	1536	2.46 PB	768	3 PB	240	96 TB	1776

1. Combined total of 2.5-in. disk drives and 1.8-in. flash cards.

Business Class configurations

The Business Class configuration is a high-density, high-performance configuration that includes standard disk enclosures and high-performance flash enclosures.

Business Class storage systems are scalable up to 16-core processors, with up to 240 flash cards, and up to 1,056 standard drives. They are optimized and configured for cost, by minimizing the number of device adapters and maximizing the number of storage enclosures attached to each storage system. They support the following storage enclosures:

- Up to 5 standard drive enclosure pairs and up to four high-performance flash enclosures in a base frame (model 961).
- Up to 7 standard drive enclosures pairs and up to 4 high-performance flash enclosures in a first expansion frame (model 96E).
- Up to 10 standard drive enclosure pairs in a second expansion frame.

The Business Class configuration uses 8 or 16 Gbps Fibre Channel host adapters that run Fibre Channel Protocol (FCP), FICON, or Fibre Channel Arbitrated Loop (FC-AL) protocol. The High Performance FICON (HPF) feature is also supported.

This configuration supports three-phase and single-phase power.

Restriction: Copy Services and I/O Priority Manager functions require a minimum of 32 GB system memory.

For more specifications, see the IBM DS8000 series specifications web site (www.ibm.com/systems/storage/disk/ds8000/specifications.html).

The following tables list the hardware components and maximum capacities that are supported for the business-class configuration, depending on the amount of memory that is available.

Table 9. Components for the Business Class configuration

Processors	System memory	Processor memory	I/O enclosures	Host adapters (8 or 4 port)	Device adapter pairs	Flash RAID adapter pairs	Standard drive enclosure pairs ^{1, 2}	Flash enclosures	Expansion frames
2-core	16 GB	8 GB	2	2 - 4	1 - 2	0	1 - 3	0	0
2-core	32 GB	16 GB	2	2 - 4	1 - 2	0 - 2	0 - 3	0 - 2	0
4-core	64 GB	32 GB	4	2 - 8	1 - 4	0 - 4	0 - 5	0 - 4	0
8-core	128 GB	64 GB	8	2 - 16	1 - 6	0 - 8	0 - 22	0 - 8	0 - 2
8-core	256 GB	128 GB	8	2 - 16	1 - 6	0 - 8	0 - 22	0 - 8	0 - 2
16-core	512 GB	256 GB	8	2 - 16	1 - 6	0 - 8	0 - 22	0 - 8	0 - 2
16-core	1,024 GB	512 GB	8	2 - 16	1 - 6	0 - 8	0 - 22	0 - 8	0 - 2

1. Standard drive enclosures are installed in pairs.
2. This configuration of the DS8870 must be populated with either one standard drive enclosure pair (feature code 1241) or one high-performance flash enclosure (feature code 1500).

Table 10. Maximum capacity for the Business Class configuration

Processors	System memory	Maximum 2.5-in. disk drives	Maximum storage capacity for 2.5-in. disk drives	Maximum 3.5-in. disk drives	Maximum storage capacity for 3.5-in. disk drives	Maximum 1.8-in. flash cards	Maximum storage capacity for 1.8-in. flash cards	Maximum total drives ¹
2-core	16 GB	144	230.4 TB	72	288 TB	N/A	N/A	144
2-core	32 GB	144	230.4 TB	72	288 TB	60	24 TB	204
4-core	64 GB	240	384 TB	120	480 TB	120	48 TB	360
8-core	128 GB	1056	1.69 PB	528	2.1 PB	240	96 TB	1296
8-core	256 GB	1056	1.69 PB	528	2.1 PB	240	96 TB	1296
16-core	512 GB	1056	1.69 PB	528	2.1 PB	240	96 TB	1296
16-core	1,024 GB	1056	1.69 PB	528	2.1 PB	240	96 TB	1296

1. Combined total of 2.5-in. disk drives and 1.8-in. flash cards.

Functional overview

The following list provides an overview of some of the features that are associated with DS8870.

Note: Some storage system functions are unavailable or are not supported in all environments. See the IBM System Storage Interoperation Center (SSIC) website (www.ibm.com/systems/support/storage/config/ssic) for the most current information on supported hosts, operating systems, adapters, and switches.

Nondisruptive and disruptive activities

DS8870 supports hardware redundancy. It is designed to support nondisruptive changes: hardware upgrades, repair, and licensed function upgrades. In addition, logical configuration changes can be made nondisruptively. For example:

- The flexibility and modularity means that expansion frames can be added and physical storage capacity can be increased within a frame without disrupting your applications.
- An increase in license scope is nondisruptive and takes effect immediately. A decrease in license scope is also nondisruptive but does not take effect until the next IML.
- Easy Tier helps keep performance optimized by periodically redistributing data to help eliminate drive hot spots that can degrade performance. This function helps balance I/O activity across the drives in an existing drive tier. It can also automatically redistribute some data to new empty drives added to a tier to help improve performance by taking advantage of the new resources. Easy Tier does this I/O activity rebalancing automatically without disrupting access to your data.

The following examples include activities that are disruptive:

- The installation of an earthquake resistance kit on a raised or nonraised floor.
- The removal of an expansion frame from the base frame.

Energy reporting

You can use DS8870 to display the following energy measurements through the DS CLI:

- Average inlet temperature in Celsius
- Total data transfer rate in MB/s
- Timestamp of the last update for values

The derived values are averaged over a 5-minute period. For more information about energy-related commands, see the commands reference.

You can also query power usage and data usage with the **showsu** command, for releases R7.2 and later. For more information, see the **showsu** description in the *Command-Line Interface User's Guide*.

National Institute of Standards and Technology (NIST) SP 800-131A security enablement

NIST SP 800-131A requires the use of cryptographic algorithms that have security strengths of 112 bits to provide data security and data integrity for secure data created in the cryptoperiod starting in 2014. The DS8870 is enabled for NIST SP 800-131A. Conformance with NIST SP 800-131A depends on the use of appropriate prerequisite management software versions and appropriate configuration of the DS8870 and other network-related entities.

Storage pool striping (rotate capacity)

Storage pool striping is supported on the DS8000 series, providing improved performance. The storage pool striping function stripes new volumes across all arrays in a pool. The striped volume layout reduces workload skew in the system without requiring manual tuning by a storage administrator. This approach can increase performance with minimal operator effort. With storage pool striping support, the system automatically performs close to highest efficiency, which requires little or no administration. The effectiveness of performance management tools is also enhanced because imbalances tend to occur as isolated problems. When performance administration is required, it is applied more precisely.

You can configure and manage storage pool striping using the DS8000 Storage Management GUI, DS CLI, and DS Open API. The rotate capacity allocation method (also referred to as rotate extents) is the default for

volumes. The rotate capacity option (storage pool striping) is designed to provide the best performance by striping volumes across arrays in the pool. Existing volumes can be reconfigured nondisruptively by using manual volume migration and volume rebalance.

The storage pool striping function is provided with the DS8000 series at no additional charge.

Performance statistics

You can use usage statistics to monitor your I/O activity. For example, you can monitor how busy the I/O ports are and use that data to help manage your SAN. For more information, see documentation about performance monitoring in the DS8000 Storage Management GUI.

Sign-on support using Lightweight Directory Access Protocol (LDAP)

The DS8000 system provides support for both unified sign-on functions (available through the DS8000 Storage Management GUI), and the ability to specify an existing Lightweight Directory Access Protocol (LDAP) server. The LDAP server can have existing users and user groups that can be used for authentication on the DS8000 system.

Setting up unified sign-on support for the DS8000 system is achieved using the Tivoli Storage Productivity Center. For more information, see the IBM Tivoli Storage Productivity Center online product documentation (www.ibm.com/support/knowledgecenter/SSNE44/).

Note: Other supported user directory servers include IBM Directory Server and Microsoft Active Directory.

Easy Tier

Easy Tier is designed to determine the appropriate tier of storage based on data access requirements and then automatically and nondisruptively move data, at the subvolume or sub-LUN level, to the appropriate tier on the DS8000 system. Easy Tier is an optional feature that offers enhanced capabilities through features such as auto-rebalancing, hot spot management, rank depopulation, and manual volume migration.

Easy Tier enables the DS8870 system to automatically balance I/O access to drives to avoid hot spots on arrays. Easy Tier can place data in the storage tier that best suits the access frequency of the data. Highly accessed data can be moved nondisruptively to a higher tier, and likewise cooler data is moved to a lower tier (for example, to Nearline drives).

Easy Tier-ing also can benefit homogeneous drive pools because it can move data away from over-utilized arrays to under-utilized arrays to eliminate hot spots and peaks in drive response times.

Z Synergy

The DS8870 storage system can work in cooperation with IBM z Systems hosts to provide the following performance enhancement functions.

- Parallel access volumes and Hyper PAV (also referred to as aliases)
- I/O Priority Manager with z/OS Workload Manager
- Extended address volumes
- High Performance FICON for IBM z Systems
- Quick initialization for IBM z Systems

Copy Services

The DS8870 storage system supports a wide variety of Copy Service

functions, including Remote Mirror, Remote Copy, and Point-in-Time functions. The key Copy Service Functions are:

- FlashCopy and FlashCopy Space Efficient
- Remote Pair FlashCopy (Preserve Mirror)
- Remote Mirror and Copy:
 - Metro Mirror
 - Global Copy
 - Global Mirror
 - Metro/Global Mirror
 - z/OS Global Mirror
 - z/OS Metro/Global Mirror

For more information on available commands, see the .

Multitenancy support (resource groups)

Resource groups provide additional policy-based limitations. Resource groups, together with the inherent volume addressing limitations, support secure partitioning of Copy Services resources between user-defined partitions. The process of specifying the appropriate limitations is performed by an administrator using resource groups functions. DS CLI support is available for resource groups functions.

Multitenancy can be supported in certain environments without the use of resource groups, provided that the following constraints are met:

- Either Copy Services functions are disabled on all DS8000 systems that share the same SAN (local and remote sites) or the landlord configures the operating system environment on all hosts (or host LPARs) attached to a SAN, which has one or more DS8000 systems, so that no tenant can issue Copy Services commands.
- The z/OS Distribute Data backup feature is disabled on all DS8000 systems in the environment (local and remote sites).
- Thin provisioned volumes (ESE or TSE) are not used on any DS8000 systems in the environment (local and remote sites).
- On zSeries systems there is only one tenant running in an LPAR, and the volume access is controlled so that a CKD base volume or alias volume is only accessible by a single tenant's LPAR or LPARs.

I/O Priority Manager

The I/O Priority Manager function can help you effectively manage quality of service levels for each application running on your system. This function aligns distinct service levels to separate workloads in the system to help maintain the efficient performance of each DS8000 volume. The I/O Priority Manager detects when a higher-priority application is hindered by a lower-priority application that is competing for the same system resources. This detection might occur when multiple applications request data from the same drives. When I/O Priority Manager encounters this situation, it delays lower-priority I/O data to assist the more critical I/O data in meeting its performance targets.

Use this function to consolidate more workloads on your system and to ensure that your system resources are aligned to match the priority of your applications.

The default setting for this feature is disabled.

Note: If the I/O Priority Manager LIC key is activated, you can enable I/O Priority Manager on the **Advanced** tab of the System settings page in the DS8000 Storage Management GUI.

Restriction of hazardous substances (RoHS)

The DS8870 system meets RoHS requirements. It conforms to the following EC directives:

- Directive 2011/65/EU of the European Parliament and of the Council of 8 June 2011 on the restriction of the use of certain hazardous substances in electrical and electronic equipment. It has been demonstrated that the requirements specified in Article 4 have been met.
- EN 50581:2012 technical documentation for the assessment of electrical and electronic products with respect to the restriction of hazardous substances.

Logical configuration

You can use either the DS8000 Storage Management GUI or the DS CLI to configure storage on the DS8000. Although the end result of storage configuration is similar, each interface has specific terminology, concepts and procedures.

Logical configuration with DS8000 Storage Management GUI

Before you configure your storage system, it is important to understand the storage concepts and sequence of system configuration.

Figure 7 illustrates the concepts of configuration.




Figure 7. Logical configuration sequence

The following concepts are used in storage configuration.

Arrays

An array, also referred to as a managed array, is a group of storage devices

that provides capacity for a pool. An array generally consists of seven or eight drives that are managed as a Redundant Array of Independent Disks (RAID)

Pools A storage pool is a collection of storage that identifies a set of storage resources. These resources provide the capacity and management requirements for arrays and volumes that have the same storage type, either fixed block (FB) or count key data (CKD).

Volumes

A volume is a fixed amount of storage on a storage device.

LSS The logical subsystem (LSS) that enables one or more host I/O interfaces to access a set of devices.

Hosts A host is the computer system that interacts with the DS8000 storage system. Hosts defined on the storage system are configured with a user-designated host type that enables DS8000 to recognize and interact with the host. Only hosts that are mapped to volumes can access those volumes.

Logical configuration of the DS8000 storage system begins with managed arrays. When you create storage pools, you assign the arrays to pools and then create volumes in the pools. FB volumes are connected through host ports to an open systems host. CKD volumes require that logical subsystems (LSSs) be created as well so that they can be accessed by an IBM z Systems host .

Pools must be created in pairs to balance the storage workload. Each pool in the pool pair is controlled by a processor node (either Node 0 or Node 1). Balancing the workload helps to prevent one node from doing most of the work and results in more efficient I/O processing, which can improve overall system performance. Both pools in the pair must be formatted for the same storage type, either FB or CKD storage. You can create multiple pool pairs to isolate workloads.

When you create a pair of pools, you can choose to automatically assign all available arrays to the pools, or assign them manually afterward. If the arrays are assigned automatically, the system balances them across both pools, so that the workload is distributed evenly across both nodes. Automatic assignment also ensures that spares and device adapter (DA) pairs are distributed equally between the pools.

If you are connecting to a z Systems host, you must create a logical subsystem (LSS) before you can create CKD volumes.

You can create a set of volumes that share characteristics, such as capacity and storage type, in a pool pair. The system automatically balances the capacity in the volume sets across both pools. If the pools are managed by Easy Tier, the capacity in the volumes is automatically distributed among the arrays. If the pools are not managed by Easy Tier, you can choose to use the rotate capacity allocation method, which stripes capacity across the arrays.

If the volumes are connecting to a z Systems host, the next steps of the configuration process are completed on the host. For more information, see .

If the volumes are connecting to an open systems host, map the volumes to the host. then add host ports to the host and map them to I/O ports on the storage system.

FB volumes can only accept I/O from the host ports of hosts that are mapped to the volumes. Host ports are zoned to communicate only with certain I/O ports on the storage system. Zoning is configured either within the storage system by using I/O port masking, or on the switch. Zoning ensures that the workload is spread properly over I/O ports and that certain workloads are isolated from one another, so that they do not interfere with each other.

The workload enters the storage system through I/O ports, which are on the host adapters. The workload is then fed into the processor nodes, where it can be cached for faster read/write access. If the workload is not cached, it is stored on the arrays in the storage enclosures.

Logical configuration with DS CLI

Before you configure your storage system with the DS CLI, it is important to understand IBM terminology for storage concepts and the storage hierarchy.

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

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

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

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

A logical volume is composed of one or more extents. A volume group specifies a set of logical volumes. By identifying different volume groups for different uses or functions (for example, SCSI target, FICON/ESCON control unit, remote mirror and copy secondary volumes, FlashCopy targets, and Copy Services), access to the set of logical volumes that are identified by the volume group can be controlled. Volume groups map hosts to volumes. Figure 8 on page 22 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. With the Quick Initialization feature for open system on CKD TSE and FB ESE or TSE 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. This

process dynamically initializes logical volumes when they are created or expanded, allowing them to be configured and placed online more quickly.

You can now specify LUN ID numbers through the graphical user interface (GUI) for volumes in a map-type volume group. Do this when you create a new volume group, add volumes to an existing volume group, or add a volume group to a new or existing host. Previously, gaps or holes in LUN ID numbers could result in a "map error" status. The Status field is eliminated from the Volume Groups main page in the GUI and the Volume Groups accessed table on the Manage Host Connections page. You can also assign host connection nicknames and host port nicknames. Host connection nicknames can be up to 28 characters, which is expanded from the previous maximum of 12. Host port nicknames can be 32 characters, which is expanded from the previous maximum of 16.




Figure 8. Logical configuration sequence

The storage management software can be used in real-time mode. 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.

RAID implementation

RAID implementation improves data storage reliability and performance.

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

collection of two or more drives presents the image of a single drive to the system. If a single device failure occurs, data can be read or regenerated from the other drives in the array.

RAID implementation provides fault-tolerant data storage by storing the data in different places on multiple drives. By placing data on multiple drives, I/O operations can overlap in a balanced way to improve the basic reliability and performance of the attached storage devices.

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

You can reconfigure RAID 5 arrays as RAID 10 arrays or vice versa.

RAID 5 overview

RAID 5 is a method of spreading volume data across multiple drives. The storage system supports RAID 5 arrays.

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

RAID 6 overview

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

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

RAID 10 overview

RAID 10 provides high availability by combining features of RAID 0 and RAID 1. The DS8000 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 disk drive to another disk drive. 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 is not required to manage parity. However, with half of the disk drives in the group used for data and the other half used to mirror that data, RAID 10 arrays have less capacity than RAID 5 arrays.

Logical subsystems

To facilitate configuration of a storage system, volumes are partitioned into groups of volumes. Each group is referred to as a logical subsystem (LSS).

As part of the storage configuration process, you can configure the maximum number of LSSs that you plan to use. The DS8000 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 volumes that have the same storage type, either count key data (CKD) for z Systems hosts or fixed block (FB) for open systems hosts.

An LSS is uniquely identified within the storage system by an identifier that consists of two hex characters (0-9 or uppercase AF) for which the volumes are associated. A fully qualified LSS is designated using the storage system identifier and the LSS identifier, such as IBM.2107-921-12FA123/1E. 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 system.

The storage system 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 system functions, including some that are associated with FB volumes, might have dependencies on LSS partitions. For example:

- The LSS partitions and their associated volume numbers must identify volumes that are specified for storage system 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 system.

If you increase storage system 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 z Systems hosts, the storage system does not process Remote Mirror and Copy functions. The FB LSS limit must be no lower than eight to support Remote Mirror and Copy functions for open-systems hosts.

Allocation methods

Allocation methods (also referred to as extent allocation methods) determine the means by which volume capacity is allocated within a pool. Allocation methods include rotate capacity, rotate volumes, and managed.

All extents of the ranks that are assigned to an extent pool are independently available for allocation to logical volumes. The extents for a LUN or volume are logically ordered, but they do not have to come from one rank and the extents do not have to be contiguous on a rank. This construction method of using fixed extents to form a logical volume in the storage system allows flexibility in the management of the logical volumes. You can delete volumes, resize volumes, and

reuse the extents of those volumes to create other volumes, different sizes. One logical volume can be deleted without affecting the other logical volumes that are defined on the same extent pool.

Because the extents are cleaned after you delete a volume, it can take some time until these extents are available for reallocation. The reformatting of the extents is a background process.

There are three allocation methods that are used by the storage system: rotate capacity (also referred to as storage pool striping), rotate volumes, and managed.

Rotate capacity allocation method

The default allocation method is rotate capacity, which is also referred to as storage pool striping. The rotate capacity allocation method is designed to provide the best performance by striping volume extents across arrays in a pool. The storage system keeps a sequence of arrays. The first array in the list is randomly picked at each power-on of the storage subsystem. The storage system tracks the array in which the last allocation started. The allocation of a first extent for the next volume starts from the next array in that sequence. The next extent for that volume is taken from the next rank in sequence, and so on. The system rotates the extents across the arrays.

If you migrate a volume with a different allocation method to a pool that has the rotate capacity allocation method, then the volume is reallocated. If you add arrays to a pool, the rotate capacity allocation method reallocates the volumes by spreading them across both existing and new arrays.

You can configure and manage this allocation method by using the DS8000 Storage Management GUI, DS CLI, and DS Open API.

Rotate volumes allocation method

Volume extents can be allocated sequentially. In this case, all extents are taken from the same array until there are enough extents for the requested volume size or the array is full, in which case the allocation continues with the next array in the pool.

If more than one volume is created in one operation, the allocation for each volume starts in another array. You might want to consider this allocation method when you prefer to manage performance manually. The workload of one volume is allocated to one array. This method makes the identification of performance bottlenecks easier; however, by putting all the volume data onto just one array, you might introduce a bottleneck, depending on your actual workload.

Managed allocation method

When a volume is managed by Easy Tier, the allocation method of the volume is referred to as managed. Easy Tier allocates the capacity in ways that might differ from both the rotate capacity and rotate volume allocation methods.

Management interfaces

You can use various IBM storage management interfaces to manage your DS8000 storage system.

These interfaces include DS8000 Storage Management GUI, DS Command-Line Interface (DS CLI), the DS Open Application Programming Interface, Tivoli Storage Productivity Center, and Tivoli Storage Productivity for Replication Manager.

Note: You can have a maximum of 256 interfaces of any type that is connected at one time.

DS8000 Storage Management GUI

Use the DS8000 Storage Management GUI interface to configure and manage storage.

DS8000 Storage Management GUI is a web-based GUI that is installed on the Hardware Management Console. You can access the DS8000 Storage Management GUI from any network-attached system by using a supported web browser. For a list of supported browsers, see “DS8000 Storage Management GUI supported web browsers” on page 29. You can also view the DS8000 Storage Management GUI by using the Element Manager in Tivoli Storage Productivity Center.

You can access the DS8000 Storage Management GUI from a browser by using the following web address, where *HMC_IP* is the IP address or host name of the HMC.
`https://HMC_IP`

If the DS8000 Storage Management GUI does not display as anticipated, clear the cache for your browser, and try to log in again.

Notes:

- If the storage system is configured for NIST SP 800-131A security conformance, a version of Java that is NIST SP 800-131A compliant must be installed on all systems that run DS8000 Storage Management GUI. For more information about security requirements, see the online product documentation about configuring your environment for NIST SP 800-131A compliance in IBM Knowledge Center (www.ibm.com/support/knowledgecenter/HW213_v7.4.0/).
- In DS8000 R7.2 and later, user names and passwords are encrypted, due to the use of the HTTPS protocol. The non-secure HTTP protocol (port 8451) is no longer available for DS8000 Storage Management GUI access.

DS command-line interface

The IBM DS command-line interface (DS CLI) can be used to create, delete, modify, and view Copy Services functions and the logical configuration of a storage unit. These tasks can be performed either interactively, in batch processes (operating system shell scripts), or in DS CLI script files. A DS CLI script file is a text file that contains one or more DS CLI commands and can be issued as a single command. DS CLI can be used to manage logical configuration, Copy Services configuration, and other functions for a storage unit, including managing security settings, querying point-in-time performance information or status of physical resources, and exporting audit logs.

The DS CLI provides a full-function set of commands to manage logical configurations and Copy Services configurations. The DS CLI can be installed on and is supported in many different environments, including the following platforms:

- AIX 5.1, 5.2, 5.3, 6.1, 7.1
- HP-UX 11.0, 11i, 11iv1, 11iv2, 11iv3
- HP Tru64 UNIX version 5.1, 5.1A

- Linux RedHat 3.0 Advanced Server (AS) and Enterprise Server (ES)
- Red Hat Enterprise Linux (RHEL) 4, RHEL 5, RHEL 6, and RHEL 7
- SuSE 8, SuSE 9, SuSE Linux Enterprise Server (SLES) 8, SLES 9, SLES 10, and SLES 11
- VMware ESX v3.0.1 Console
- Novell NetWare 6.5
- IBM System i® i5/OS 5.4, 6.1, 7.1
- OpenVMS 7.3-1 (or newer, Alpha processor only)
- Oracle Solaris 7, 8, and 9
- Microsoft Windows Server 2000, 2003, 2008, 2012, Windows Datacenter, Windows XP, Windows Vista, Windows 7, 8

Note: If the storage system is configured for NIST SP 800-131A security conformance, a version of Java that is NIST SP 800-131A compliant must be installed on all systems that run DS CLI. For more information about security requirements, see documentation about configuring your environment for NIST SP 800-131A compliance in IBM Knowledge Center (www.ibm.com/support/knowledgecenter/HW213_v7.4.0/).

DS Open Application Programming Interface

The DS Open Application Programming Interface (API) is a nonproprietary storage management client application that supports routine LUN management activities. Activities that are supported include: LUN creation, mapping and masking, and the creation or deletion of RAID 5, RAID 6, and RAID 10 volume spaces.

The DS Open API supports these activities through the use of the Storage Management Initiative Specification (SMI-S), as defined by the Storage Networking Industry Association (SNIA).

The DS Open API helps integrate configuration management support into storage resource management (SRM) applications, which help you to use existing SRM applications and infrastructures. The DS Open API can also be used to automate 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 DS8000 Storage Management GUI web-based interface and the DS command-line interface.

Note: The DS Open API supports the storage system and is an embedded component.

You can implement the DS Open API without using a separate middleware application. For example, you can implement it with the IBM Common Information Model (CIM) agent, which 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 is used by storage management applications to communicate with a storage unit.

IBM Storage Mobile Dashboard

IBM Storage Mobile Dashboard is a free application that provides basic monitoring capabilities for IBM storage systems. You can securely check the health and performance status of your IBM DS8000 series storage system by viewing events and performance metrics.

To install IBM Storage Mobile Dashboard on an iOS device, open the App Store app and search for “IBM Storage Mobile Dashboard.”

Tivoli Storage Productivity Center

The Tivoli Storage Productivity Center is an integrated software solution that can help you improve and centralize the management of your storage environment through the integration of products. With the Tivoli Storage Productivity Center (TPC), it is possible to manage and fully configure multiple DS8000 systems from a single point of control.

Note: The installation of TPC is not required for the operation of DS8870. However, it is recommended. TPC can be ordered and installed as a software product on a variety of servers and operating systems. When installing TPC, ensure that the selected version supports your use of the latest system functions. Optionally, you can order a server in which TPC is preinstalled.

DS8000 Storage Management GUI is a web-based GUI that is installed on the Hardware Management Console. You can view DS8000 Storage Management GUI from any network-attached system using a supported web browser. You can also view DS8000 Storage Management GUI by using the Element Manager in Tivoli Storage Productivity Center.

Tivoli Storage Productivity Center simplifies storage management by providing the following benefits:

- Centralizing the management of heterogeneous storage network resources with IBM storage management software
- Providing greater synergy between storage management software and IBM storage devices
- Reducing the number of servers that are required to manage your software infrastructure
- Migrating from basic device management to storage management applications that provide higher-level functions

With the help of agents, Tivoli Storage Productivity Center discovers the devices to which it is configured. It then can start an element manager that is specific to each discovered device, and gather events and data for reports about storage management.

For more information, see IBM Tivoli Storage Productivity Center online product documentation (www.ibm.com/support/knowledgecenter/SSNE44/).

Tivoli Storage Productivity Center for Replication

Tivoli Storage Productivity Center for Replication facilitates the use and management of Copy Services functions such as the remote mirror and copy functions (Metro Mirror and Global Mirror) and the point-in-time function (FlashCopy).

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

Tivoli Storage Productivity Center for Replication for FlashCopy, Metro Mirror, and Global Mirror support provides automation of administration and configuration of

these services, operational control (starting, suspending, resuming), Copy Services tasks, and monitoring and managing of copy sessions.

Tivoli Storage Productivity Center for Replication supports Multiple Target PPRC for data replication from a single primary site to two secondary sites simultaneously.

Tivoli Storage Productivity Center for Replication is part of Tivoli Storage Productivity Center V5 software program. If you are licensed for Copy Services functions, you can use Tivoli Storage Productivity Center to manage your Copy Services environment.

Notes:

1. You can connect to a storage system on a hardware management console (HMC) by using the Tivoli Storage Productivity Center for Replication.
2. The use of Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6) are both supported through the HMC ports.

For more information, see the IBM Tivoli Storage Productivity Center online product documentation (www.ibm.com/support/knowledgecenter/SSNE44/).

DS8000 Storage Management GUI supported web browsers

To access the DS8000 Storage Management GUI, you must ensure that your web browser is supported and has the appropriate settings enabled.

The DS8000 Storage Management GUI supports the following web browsers:

Table 11. Supported browsers according to DS8000 version

DS8000 version	Supported browsers
7.5	Mozilla Firefox 35 Mozilla Firefox Extended Support Release (ESR) 31 Microsoft Internet Explorer 11 Google Chrome 41
7.4	Mozilla Firefox 30 Mozilla Firefox Extended Support Release (ESR) 24 Microsoft Internet Explorer 10 and 11 Google Chrome 36
7.2 and 7.3	Microsoft Internet Explorer 9 Mozilla Firefox 17 ESR
7.0 and 7.1	Microsoft Internet Explorer 9 Mozilla Firefox 10 ESR

IBM supports higher versions of the browsers as long as the vendors do not remove or disable functionality that the product relies upon. For browser levels higher than the versions that are certified with the product, customer support accepts usage-related and defect-related service requests. As with operating system and virtualization environments, if the support center cannot re-create the issue in our lab, we might ask the client to re-create the problem on a certified browser version to determine whether a product defect exists. Defects are not accepted for cosmetic differences between browsers or browser versions that do not affect the functional behavior of the product. If a problem is identified in the product, defects are accepted. If a problem is identified with the browser, IBM might investigate potential solutions or workaround that the client can implement until a permanent solution becomes available.

Enabling TLS 1.2 support

If the security requirements for your storage system require conformance with NIST SP 800-131A, enable transport layer security (TLS) 1.2 on web browsers that use SSL/TLS to access the DS8000 Storage Management GUI. See your web browser documentation for instructions on enabling TLS 1.2. For Internet Explorer, complete the following steps to enable TLS 1.2.

1. On the **Tools** menu, click **Internet Options**.
2. On the **Advanced** tab, under **Settings**, select **Use TLS 1.2**.

Note: Firefox, Release 24 and later, supports TLS 1.2. However, you must configure Firefox to enable TLS 1.2 support.

For more information about security requirements, see .

Selecting browser security settings

You must select the appropriate web browser security settings to access the DS8000 Storage Management GUI. In Internet Explorer, use the following steps.

1. On the **Tools** menu, click **Internet Options**.
2. On the **Security** tab, select **Internet** and click **Custom level**.
3. Scroll to **Miscellaneous**, and select **Allow META REFRESH**.
4. Scroll to **Scripting**, and select **Active scripting**.

Configuring Internet Explorer to access the DS8000 Storage Management GUI

If DS8000 Storage Management GUI is accessed through the Tivoli Storage Productivity Center with Internet Explorer, use the following steps to properly configure the web browser.

1. Configure Internet Explorer to disable the **Pop-up Blocker**. From the Internet Explorer toolbar, click **Tools > Pop-up Blocker > Turn Off Pop-up Blocker**.

Note: If a message indicates that content is blocked without a signed by a valid security certificate, click the **Information Bar** at the top and select **Show blocked content**.

2. Add the IP address of the DS8000 Hardware Management Console (HMC) to the Internet Explorer list of trusted sites.
 - a. From the Internet Explorer toolbar, click **Tools > Internet options**.
 - b. On the Security tab, click the **Trusted sites** icon and then click **Sites**.
 - c. In the **Add this web site to the zone** field, type the IP address of the DS8000 Hardware Management Console.
 - d. Click **Add** to add the IP address to the appropriate field, and then click **OK**.
 - e. Click **OK** to exit Internet Options, and then close Internet Explorer.

Chapter 2. Hardware features

Use this information to assist you with planning, ordering, and managing your DS8000 series.

The following table lists feature codes that are used to order hardware features for DS8000 series.

Table 12. Feature codes for hardware features

Feature code	Feature	Description
0200	Shipping weight reduction	Maximum shipping weight of any storage system base model or expansion model does not exceed 909 kg (2000 lb) each. Packaging adds 120 kg (265 lb).
0400	BSMI certification documents	Required when the storage system model is shipped to Taiwan.
1000	Remote zSeries power control	An optional feature that is used to control power on/off sequence from a z Systems server.
1051	Battery service modules	Required for each frame.
1055	Extended power line disturbance	An optional feature that is used to protect the storage system from a power-line disturbance for up to 50 seconds.
1061	Single-phase power cord, 200 - 240 V, 60 A, 3-pin connector	HBL360C6W, Pin and Sleeve Connector, IEC 309, 2P3W HBL360R6W, AC Receptacle, IEC 60309, 2P3W
1068	Single-phase power cord, 200 - 240 V, 63 A, no connector	Inline Connector: not applicable Receptacle: not applicable
1072	Top exit SPP, 200-240V, 60 A, 3-pin connector	6 mm ²
1073	Top exit SPP, 200-240V, 63 A, no connector	6 mm ²
1080	Three-phase power cord, high voltage (five-wire 3P+N+G) 380-415V (nominal), 30 A, IEC 60309, 5-pin customer connector	HBL530C6V02, Pin and Sleeve Connector, IEC 60309, 4P5W HBL530R6V02, AC Receptacle, IEC 603309, 4P5W
1081	Three-phase high voltage (five-wire 3P+N+G), 380-415V, 32 A, no customer connector provided	Inline Connector: not applicable Receptacle: not applicable
1082	Three-phase power cord, low voltage, 200 -240 V, 60A, 4-pin connector	HBL460C9W, Pin and Sleeve Connector, IEC 309, 3P4W HBL460R9W, AC Receptacle, IEC 60309, 3P4W

Table 12. Feature codes for hardware features (continued)

Feature code	Feature	Description
1083	Top exit three-phase high voltage (five-wire 3P+N+G), 380-415V, 32 A, no customer connector provided	Inline Connector: not applicable Receptacle: not applicable
1084	Top exit three-phase low voltage (four-wire 3P+G), 200-240V, 60 A, IEC 60309 4-pin customer connector	HBL460C9W, Pin and Sleeve Connector, IEC 309, 3P4W HBL460R9W, AC Receptacle, IEC 60309, 3P4W
1085	Top exit three-phase high voltage (five-wire 3P+N+G), 380-415V (nominal), 30 A, IEC 60309 5-pin customer connector	HBL530C6V02, Pin and Sleeve Connector, IEC 603309, 4P5W HBL530R6V02, AC Receptacle, IEC 603309, 4P5W
1095	Top exit power cord	For U.S., Canada, Latin America, and Asia Pacific (available for only Models 951 and 95E)
1096	Top exit power cord	For Africa (available for only Models 951 and 95E)
1097	Top exit power cord	For U.S., Canada, and Japan (available for only Models 951 and 95E)
1101	Universal ladder for top exit cable access	Required for all top-exit power cords and top-exit tailgate.
1120	Internal management console (notebook)	A required feature that is installed in the 961 frame
1130	External management console (notebook)	An optional feature that can be installed in an external IBM or a non-IBM rack
1170	Management-console power cord standard rack	
1171	Management-console power cord group 1	Only for United States, Canada, Bahamas, Barbados, Bermuda, Bolivia, Brazil, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Guyana, Honduras, Jamaica, Japan, Japan (PDS), Mexico, Netherlands Antilles, Panama, Philippines, Saudi Arabia, Suriname, Taiwan, Trinidad, Venezuela
1172	Management-console power cord group 2	Only for Brazil
1241	Standard drive enclosure pair	For 2.5-inch disk drives
1242	Standard drive enclosure	For 2.5-inch disk drives
1244	Standard drive enclosure pair	For 3.5-inch disk drives
1245	Standard drive enclosure	For 400 GB flash drives
1246	Drive cable group 1	Connects the disk drives to the device adapters within the same base model 961.

Table 12. Feature codes for hardware features (continued)

Feature code	Feature	Description
1247	Drive cable group 2	(Enterprise-class) Connects the disk drives to the device adapters in the first expansion model 96E (Business-class) Connects the drives from the first expansion model 96E to the base model 961.
1248	Drive cable group 4	Connects the disk drives from a second expansion model 96E to the base model 961 and first expansion model 96E.
1249	Drive cable group 5	(Enterprise-class) Connects the disk drives from a third expansion model 96E to a second expansion model 96E. (Business-class) Not applicable
1250	Drive cable group 1	(DS8870 business-class only) Connects the disk drives from a third expansion model 96E to the second expansion model 96E.
1255	Standard drive enclosure	For 200 GB flash drives
1256	Standard drive enclosure	For 800 GB flash drives
1257	1.6 TB SSD enclosure indicator	
1301	I/O enclosure pair	
1320	PCIe cable group 1	Connects device and host adapters in an I/O enclosure pair to the processor.
1321	PCIe cable group 2	Connects device and host adapters in I/O enclosure pairs to the processor.
1322	PCIe cable group 3	Connects device and host adapters in an I/O enclosure pair to the processor.
1400	Top-exit bracket for Fibre cable	
1410	Fibre Channel cable	40 m (131 ft), 50 micron OM3 or higher, multimode
1411	Fibre Channel cable	31 m (102 ft), 50 micron OM3 or higher, multimode
1412	Fibre Channel cable	2 m (6.5 ft) , 50 micron OM3 or higher, multimode
1420	Fibre Channel cable	31 m (102 ft), 9 micron OM3 or higher, single mode
1421	Fibre Channel cable	31 m (102 ft), 9 micron OM3 or higher, single mode
1422	Fibre Channel cable	2 m (6.5 ft), 9 micron OM3 or higher, single mode
1500	High performance flash enclosure	For flash cards
1506	400 GB 1.8-inch flash cards set	Flash card set A (16 cards)

Table 12. Feature codes for hardware features (continued)

Feature code	Feature	Description
1508	400 GB 1.8-inch flash cards set	Flash card set B (14 cards)
1596	400 GB 1.8-inch flash cards set	Flash card set A (16 cards)
1598	400 GB 1.8-inch flash cards set	Flash card set B (14 cards)
1599	Flash enclosure filler set	Includes 14 fillers
1735	DS8000 Licensed Machine Code R7.4	Microcode bundle 87.x.xx.x
1736	DS8000 Licensed Machine Code R7.5	Microcode bundle 87.x.xx.x
1760	External Security Key Life-cycle Manager (SKLM) isolated-key appliance	Model 961
1761	External SKLM isolated-key appliance	Model AP1 single server configuration
1762	Secondary external SKLM isolated-key appliance	Model AP1 dual server configuration
1906	Earthquake resistance kit	One of two versions of this kit might be used. Ensure that you use the correct template for floor preparation.
2997	Disk enclosure filler set	For 3.5-in. DDMs; includes eight fillers
2998	Disk enclosure filler set	For 2.5-in. DDMs; includes eight fillers
2999	Disk enclosure filler set	For 2.5-in. DDMs; includes 16 fillers
3053	Device adapter pair	4-port, 8 GB
3054	Flash enclosure adapter pair	Required for feature code 1500
3153	Fibre Channel host-adapter cable	4-port, 8 Gbps shortwave FCP and FICON host adapter PCIe
3157	Fibre Channel host-adapter cable	8-port, 8 Gbps shortwave FCP and FICON host adapter PCIe
3253	Fibre Channel host-adapter cable	4-port, 8 Gbps longwave FCP and FICON host adapter PCIe
3257	Fibre Channel host-adapter cable	8-port, 8 Gbps longwave FCP and FICON host adapter PCIe
3353	Fibre Channel host-adapter cable	4-port, 16 Gbps shortwave FCP and FICON host adapter PCIe
3453	Fibre Channel host-adapter cable	4-port, 16 Gbps longwave FCP and FICON host adapter PCIe
4311	16 GB system memory	(2-core)
4312	32 GB system memory	(2-core)
4313	64 GB system memory	(4-core)
4314	128 GB system memory	(8-core)
4315	256 GB system memory	(8-core)
4316	512 GB system memory	(16-core)
4317	1 TB system memory	(16-core)
4401	2-core POWER7 processors	Requires feature code 4311 or 4312

Table 12. Feature codes for hardware features (continued)

Feature code	Feature	Description
4402	4-core POWER7 processors	Requires feature code 4313
4403	8-core POWER7 processors	Requires feature code 4314 or 4315
4404	16-core POWER7 processors	Requires feature code 4316 or 4317
4411	2-core POWER7+ processors	Requires feature code 4311 or 4312
4412	4-core POWER7+ processors	Requires feature code 4313
4413	8-core POWER7+ processors	Requires feature code 4314 or 4315
4414	16-core POWER7+ processors	Requires feature code 4316 or 4317
5108	146 GB 15 K FDE disk-drive set	SAS
5209	146 GB 15 K FDE CoD disk-drive set	SAS
5308	300 GB 15 K FDE disk-drive set	SAS
5309	300 GB 18 K FDE CoD disk-drive set	SAS
5618	600 GB 15 K FDE disk-drive set	SAS
5619	600 GB 15 K FDE CoD disk-drive set	SAS
5708	600 GB 10K FDE disk-drive set	SAS
5709	600 GB 10K FDE CoD disk-drive set	SAS
5758	900 GB 10K FDE disk-drive set	SAS
5759	900 GB 10K FDE CoD disk-drive set	SAS
5768	1.2 TB 10K FDE disk-drive set	SAS
5769	1.2 TB 10K FDE CoD disk-drive set	SAS
5858	3 TB 7.2 K FDE half disk-drive set	SAS
5859	3 TB 7.2 K FDE CoD disk-drive drive set	SAS
5868	4 TB 7.2 K FDE disk-drive set	SAS
5869	4 TB 7.2 K FDE CoD disk-drive set	SAS
6058	200 GB FDE flash-drive set	SAS
6156	400 GB FDE half flash-drive set	SAS
6158	400 GB FDE flash-drive set	SAS
6258	800 GB FDE flash-drive set	SAS
6358	1.6 TB SSD FDE drive set	SAS

Storage complexes

A storage complex is a set of storage units that are managed by management console units.

You can associate one or two management console units with a storage complex. Each storage complex must use at least one of the internal management console units in one of the storage units. You can add a second management console for redundancy. The second storage management console can be either one of the internal management console units in a storage unit or an external management console.

Management console

The management console supports storage system hardware and firmware installation and maintenance activities.

The *management console* is a dedicated notebook that is physically located (installed) inside your storage system, and can automatically monitor the state of your system, and notify you and IBM when service is required.

One management console in the storage system is internal. To provide continuous availability of your access to the management-console functions, use an additional management console, especially for storage environments that use encryption. For more information, see Chapter 9, "Planning for security," on page 195.

An additional management console can be provided in two ways:

External -

The external management console is installed in the customer-provided rack. This option uses the same hardware as the internal management console.

Note: The external management console must be within 50 feet of the base frame.

Internal -

The internal management console from each of two separate storage facilities can be "cross-coupled." Plan for this configuration to be accomplished during the initial installation of the two storage facilities to avoid more power cycling. (Combining two previously installed storage facilities into the cross-coupled configuration later, requires a power cycle of the second storage facility.) Ensure that you maintain the same machine code level for all storage facilities in the cross-coupled configuration.

RAID implementation

RAID implementation improves data storage reliability and performance.

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

RAID implementation provides fault-tolerant data storage by storing the data in different places on multiple drives. By placing data on multiple drives, I/O operations can overlap in a balanced way to improve the basic reliability and performance of the attached storage devices.

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

You can reconfigure RAID 5 arrays as RAID 10 arrays or vice versa.

RAID 5 overview

RAID 5 is a method of spreading volume data across multiple drives. The storage system supports RAID 5 arrays.

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

RAID 6 overview

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

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

RAID 10 overview

RAID 10 provides high availability by combining features of RAID 0 and RAID 1. The DS8000 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 disk drive to another disk drive. 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 is not required to manage parity. However, with half of the disk drives in the group used for data and the other half used to mirror that data, RAID 10 arrays have less capacity than RAID 5 arrays.

Hardware specifics

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

Contributors to the high degree of availability and reliability include the structure of the storage unit, the host systems that are supported, and the memory and speed of the processors.

Storage system structure

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

Storage system

The storage unit contains a storage server and one or more pairs of storage enclosures that are packaged in one or more racks with associated power supplies, batteries, and cooling.

Storage server

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

Processor complex

The processor complex controls and manages the storage server functions in the storage system. The two processor complexes form a redundant pair such that if either processor complex fails, the remaining processor complex controls and manages all storage server functions.

Rack power control card

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

Disk drives

DS8870 provides you with a choice of drives.

The following drives are available:

- 1.8-inch flash cards with FDE:
 - 400 GB
- 2.5-inch flash drives with FDE:
 - 200 GB
 - 400 GB
 - 800 GB
 - 1.6 TB
- 2.5-inch disk drives with Full Disk Encryption (FDE) and Standby CoD:
 - 146 GB, 15 K RPM
 - 300 GB, 15 K RPM
 - 600 GB, 10 K RPM
 - 600 GB, 15 K RPM
 - 900 GB, 10 K RPM
 - 1.2 TB 10 K RPM
- 3.5-inch disk drives with FDE and Standby CoD:
 - 3 TB, 7.2 K RPM
 - 4 TB, 7.2 K RPM

Drive maintenance policy

The DS8000 internal maintenance functions use an Enhanced Sparing process that delays a service call for drive replacement if there are sufficient spare drives. All drive repairs are managed according to Enhanced Sparing rules.

A minimum of two spare drives are allocated in a device adapter loop, and up to four spares when the number of drives reaches 32 in a specific loop. Internal maintenance functions continuously monitor and report (by using the call home feature) to IBM when the number of drives in a spare pool reaches a preset threshold. This design ensures continuous availability of devices while it protects data and minimizing any service disruptions.

It is not recommended to replace a drive unless an error is generated indicating that service is needed.

Host attachment overview

The DS8000 series provides various host attachments so that you can consolidate storage capacity and workloads for open-systems hosts and z Systems.

The DS8000 series provides extensive connectivity using Fibre Channel adapters across a broad range of server environments.

Host adapter intermix support

Both 4-port and 8-port host adapters (HAs) are available in DS8870, but like the DS8800, it can use the same 8 Gbps, 4-port HA for improved performance.

The following table shows the 4-port or 8-port host adapter plug order for by using a 4-port or 8-port HA configuration. The host adapter installation order for the second frame (I/O enclosures 5 through 8) is the same for the first four I/O enclosures.

Table 13. Plug order for 4- and 8-port HA slots (8 GB) for two and four I/O enclosures

I/O enclosures	Slot number					
	C1	C2	C3	C4	C5	C6
For two I/O enclosures in an Enterprise Class or Business Class configuration						
Top I/O enclosure 1	–	–		–	–	
Bottom I/O enclosure 3	3	X		1	X	
Top I/O enclosure 2	–	–		–	–	
Bottom I/O enclosure 4	2	X		4	X	
For four I/O enclosures in an Enterprise Class or Business Class configuration						
Top I/O enclosure 1	7	X		3	X	
Bottom I/O enclosure 3	5	X		1	X	
Top I/O enclosure 2	4	X		8	X	
Bottom I/O enclosure 4	2	X		6	X	
For eight I/O enclosures in an All Flash configuration						
Top I/O enclosure 4	15	X		7	X	
Bottom I/O enclosure 6	11	X		3	X	
Top I/O enclosure 5	8	X		16	X	
Bottom I/O enclosure 7	4	X		12	X	
Top I/O enclosure 0	13	X		5	X	

Table 13. Plug order for 4- and 8-port HA slots (8 GB) for two and four I/O enclosures (continued)

I/O enclosures	Slot number					
	C1	C2	C3	C4	C5	C6
Bottom I/O enclosure 2	9	X		1	X	
Top I/O enclosure 1	6	X		14	X	
Bottom I/O enclosure 3	2	X		10	X	

Open-systems host attachment with Fibre Channel adapters

You can attach a DS8000 series to an open-systems host with Fibre Channel adapters.

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

The IBM DS8000 series supports SAN connections of up to 8 Gbps with 8 Gbps host adapters. The IBM DS8000 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 z Systems hosts overview

The DS8000 series can be attached to FICON attached z Systems host operating systems under specified adapter configurations.

Each storage system Fibre Channel adapter has either four ports or eight ports, depending on the adapter speed. Each port has a unique worldwide port name (WWPN). You can configure the port to operate with the FICON upper-layer protocol. For FICON, the Fibre Channel port supports connections to a maximum of 509 FICON hosts. On FICON, the Fibre Channel adapter can operate with fabric or point-to-point topologies.

With Fibre Channel adapters that are configured for FICON, the storage system provides the following configurations:

- Either fabric or point-to-point topologies
- A maximum of eight 8-port host adapters on the DS8870, Model 961 (2-core) and a maximum of 16 8-port host adapters on the DS8870, Model 96E (4-core), which equates to a maximum of 128 Fibre Channel ports

Note: If 4-port 16 Gb host adapters are used in combination with the 8-port adapters, the maximum number of ports varies from 64 to 128, depending on the combination of adapters.

- A maximum of 509 logins per Fibre Channel port.
- A maximum of 8192 logins per storage system.

- A maximum of 1280 logical paths on each Fibre Channel port.
- Access to all 255 control-unit images (8000 CKD devices) over each FICON port.
- A maximum of 512 logical paths per control unit image.

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

The storage system supports the following operating systems for z Systems hosts:

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

For the most current information on supported hosts, operating systems, adapters, and switches, go to the IBM System Storage Interoperation Center (SSIC) website (www.ibm.com/systems/support/storage/config/ssic).

Processor memory

The DS8000 series offers a number of configuration options for processor memory.

2-core configuration

Enterprise Class and Business Class offer 16 or 32 GB of processor memory.

4-core configuration

Enterprise Class and Business Class offer 64 GB of processor memory.

8-core configuration

Enterprise Class and Business Class offer 128 or 256 GB of processor memory. DS8870 All Flash offers 256 GB of processor memory.

16-core configuration

Enterprise Class and Business Class offer 512 or 1024 GB of processor memory. DS8870 All Flash offers 512 or 1024 GB of processor memory.

The nonvolatile storage (NVS) scales with the selected processor memory size, which can also help optimize performance. The NVS is typically 1/32 of the installed memory.

Note: The minimum NVS is 1 GB.

Subsystem device driver for open-systems

The IBM Multipath Subsystem Device Driver (SDD) supports open-systems hosts.

All storage system models include the IBM System Storage Multipath Subsystem Device Driver (SDD). The SDD provides load balancing and enhanced data availability capability in configurations with more than one I/O path between the host server and the storage system. Load balancing can reduce or eliminate I/O

bottlenecks that occur when many I/O operations are directed to common devices by using the same I/O path. The SDD can eliminate the single point of failure by automatically rerouting I/O operations when a path failure occurs.

I/O load balancing

You can maximize the performance of an application by spreading the I/O load across processor nodes, arrays, and device adapters in the storage system.

During an attempt to balance the load within the storage system, 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 drive groups. Use as many RAID drive groups as possible for the critical applications. Most performance bottlenecks occur because a few drive are overloaded. Spreading an application across multiple RAID drive groups ensures that as many drives as possible are available. This is extremely important for open-system environments where cache-hit ratios are usually low.
- Activity to the nodes. When selecting RAID drive groups for a critical application, spread them across separate nodes. Because each node has separate memory buses and cache memory, this maximizes the use of those resources.
- Activity to the device adapters. When selecting RAID drive groups within a cluster for a critical application, spread them across separate device adapters.
- Activity to the Fibre Channel ports. Use the IBM 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 Multipath Subsystem Device Driver User's Guide* (<http://www-01.ibm.com/support/docview.wss?uid=ssg1S7000303>). This document also describes the product engineering tool, the ESSUTIL tool, which is supported in the **pcmpath** commands and the **datapath** commands.

Storage consolidation

When you use a storage system, you can consolidate data and workloads from different types of independent hosts into a single shared resource.

You can mix production and test servers in an open systems environment or mix open systems and z Systems hosts. In this type of environment, servers rarely, if ever, contend for the same resource.

Although sharing resources in the storage system has advantages for storage administration and resource sharing, there are more implications for workload planning. The benefit of sharing is that a larger resource pool (for example, 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 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 drive groups. Data for open systems or z Systems hosts is automatically placed on separate arrays, which reduce the contention for drive use.

- On separate device adapters.
- In separate processor nodes, which isolate use of memory buses, microprocessors, and cache resources. Before you decide, verify that the isolation of your data to a single node provides adequate data access performance for your application.

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.

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. A physical disk contains 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.

T10 DIF support

American National Standards Institute (ANSI) T10 Data Integrity Field (DIF) standard is supported on IBM z Systems for SCSI end-to-end data protection on fixed block (FB) LUN volumes. This support applies to the IBM Storage DS8870 unit (models 961 and 96E). IBM z Systems support applies to FCP channels only.

IBM z Systems provides added end-to-end data protection between the operating system and the DS8870 unit. This support adds protection information consisting of CRC (Cyclic Redundancy Checking), LBA (Logical Block Address), and host application tags to each sector of FB data on a logical volume.

Data protection using the T10 Data Integrity Field (DIF) on FB volumes includes the following features:

- Ability to convert logical volume formats between standard and protected formats supported through PPRC between standard and protected volumes
- Support for earlier versions of T10-protected volumes on the DS8870 with non T10 DIF-capable hosts
- Allows end-to-end checking at the application level of data stored on FB disks
- Additional metadata stored by the storage facility image (SFI) allows host adapter-level end-to-end checking data to be stored on FB disks independently of whether the host uses the DIF format.

Notes:

- This feature requires changes in the I/O stack to take advantage of all the capabilities the protection offers.
- T10 DIF volumes can be used by any type of Open host with the exception of iSeries, but active protection is supported only for Linux on IBM z Systems or AIX on IBM Power Systems. The protection can only be active if the host server has T10 DIF enabled.

- T10 DIF volumes can accept SCSI I/O of either T10 DIF or standard type, but if the FB volume type is standard, then only standard SCSI I/O is accepted.

Logical volumes

A logical volume is the storage medium that is associated with a logical disk. It typically resides on two 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.

Allocation, deletion, and modification of volumes

Extent allocation methods (namely, rotate volumes and pool striping) determine the means by which actions are completed on storage system volumes.

All extents of the ranks assigned to an extent pool are independently available for allocation to logical volumes. The extents for a LUN or volume are logically ordered, but they do not have to come from one rank and the extents do not have to be contiguous on a rank. This construction method of using fixed extents to form a logical volume in the storage system allows flexibility in the management of the logical volumes. You can delete volumes, resize volumes, and reuse the extents of those volumes to create other volumes, different sizes. One logical volume can be deleted without affecting the other logical volumes defined on the same extent pool.

Because the extents are cleaned after you delete a volume, it can take some time until these extents are available for reallocation. The reformatting of the extents is a background process.

There are two extent allocation methods used by the storage system: rotate volumes and storage pool striping (rotate extents).

Storage pool striping: extent rotation

The default storage allocation method is storage pool striping. The extents of a volume can be striped across several ranks. The storage system keeps a sequence of ranks. The first rank in the list is randomly picked at each power on of the storage subsystem. The storage system tracks the rank in which the last allocation started. The allocation of a first extent for the next volume starts from the next rank in that sequence. The next extent for that volume is taken from the next rank in sequence, and so on. The system rotates the extents across the ranks.

If you migrate an existing non-striped volume to the same extent pool with a rotate extents allocation method, then the volume is "reorganized." If you add more ranks to an existing extent pool, then the "reorganizing" existing striped volumes spreads them across both existing and new ranks.

You can configure and manage storage pool striping using the DS Storage Manager, DS CLI, and DS Open API. The default of the extent allocation method (EAM) option that is allocated to a logical volume is now rotate extents. The rotate extents option is designed to provide the best performance by striping volume extents across ranks in extent pool.

Managed EAM: Once a volume is managed by Easy Tier, the EAM of the volume is changed to managed EAM, which can result in placement of the extents differing from the rotate volume and rotate extent rules. The EAM only changes when a volume is manually migrated to a non-managed pool.

Rotate volumes allocation method

Extents can be allocated sequentially. In this case, all extents are taken from the same rank until there are enough extents for the requested volume size or the rank is full, in which case the allocation continues with the next rank in the extent pool.

If more than one volume is created in one operation, the allocation for each volume starts in another rank. When allocating several volumes, rotate through the ranks. You might want to consider this allocation method when you prefer to manage performance manually. The workload of one volume is going to one rank. This method makes the identification of performance bottlenecks easier; however, by putting all the volumes data onto just one rank, you might introduce a bottleneck, depending on your actual workload.

LUN calculation

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

In the storage system, physical storage capacities are 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.

On open volumes with 512 byte blocks (including T10-protected volumes), you can specify an exact block count to create a LUN. You can specify a standard LUN size (which is expressed as an exact number of binary GBs (2^{30})) or you can specify an ESS volume size (which is expressed in decimal GBs (10^9) accurate to 0.1 GB). The unit of storage allocation for open volumes is fixed block one extent. The extent size for open volumes is exactly 1 GB (2^{30}). Any logical volume that is not an exact multiple of 1 GB does not use all the capacity in the last extent that is allocated to the logical volume. Supported block counts are from 1 to 4 194 304 blocks (2 binary TB) in increments of one block. Supported sizes are from 1 to 2048 GB (2 binary TB) in increments of 1 GB. The supported ESS LUN sizes are limited to the exact sizes that are specified from 0.1 to 982.2 GB (decimal) in increments of 0.1 GB and are rounded up to the next larger 32 K byte boundary. The ESS LUN sizes do not result in standard LUN sizes. Therefore, they can waste capacity. However, the unused capacity is less than one full extent. ESS LUN sizes are typically used when volumes must be copied between the storage system and ESS.

On open volumes with 520 byte blocks, you can select one of the supported LUN sizes that are used on IBM i processors to create a LUN. The operating system uses 8 of the bytes in each block. This leaves 512 bytes per block for your data. Variable volume sizes are also supported.

Table 14 shows the disk capacity for the protected and unprotected models. Logically unprotecting a storage LUN allows the IBM i host to start system level mirror protection on the LUN. The IBM i system level mirror protection allows normal system operations to continue running in the event of a failure in an HBA, fabric, connection, or LUN on one of the LUNs in the mirror pair.

Note: On IBM i, logical volume sizes in the range 17.5 GB to 141.1 GB are supported as load source units. Logical volumes smaller than 17.5 GB or larger than 141.1 GB cannot be used as load source units.

Table 14. Capacity and models of disk volumes for IBM i hosts running IBM i operating system

Size	Type	Protected model	Unprotected model
8.5 GB	242x	A01	A81
17.5 GB	242x	A02	A82
35.1 GB	242x	A05	A85
70.5 GB	242x	A04	A84
141.1 GB	242x	A06	A86
282.2 GB	242x	A07	A87
1 GB to 2000 GB	242x	099	050

On CKD volumes, you can specify an exact cylinder count or a standard volume size to create a LUN. The standard volume size is expressed as an exact number of Mod 1 equivalents (which is 1113 cylinders). The unit of storage allocation for CKD volumes is one CKD extent. The extent size for CKD volume is exactly a Mod 1 equivalent (which is 1113 cylinders). Any logical volume that is not an exact multiple of 1113 cylinders (1 extent) does not use all the capacity in the last extent that is allocated to the logical volume. For CKD volumes that are created with 3380 track formats, the number of cylinders (or extents) is limited to either 2226 (1 extent) or 3339 (2 extents). For CKD volumes that are created with 3390 track formats, you can specify the number of cylinders in the range of 1 - 65520 (x'0001' - x'FFF0') in increments of one cylinder, or as an integral multiple of 1113 cylinders between 65,667 - 262,668 (x'10083' - x'4020C') cylinders (59 - 236 Mod1 equivalents). Alternatively, for 3390 track formats, you can specify Mod 1 equivalents in the range of 1-236.

Extended address volumes for CKD

Count key data (CKD) volumes now support the additional capacity of 1 TB. The 1 TB capacity is an increase in volume size from the previous 223 GB.

This increased volume capacity is referred to as extended address volumes (EAV) and is supported by the 3390 Model A. Use a maximum size volume of up to 1,182,006 cylinders for the IBM zOS. This support is available to you for the z/OS version 12.1, and later.

You can create a 1 TB IBM z Systems CKD volume on the DS8870.

A z Systems CKD volume is composed of one or more extents from a CKD extent pool. CKD extents are 1113 cylinders in size. When you define a z Systems CKD volume, you must specify the number of cylinders that you want for the volume. The storage system and the zOS have limits for the CKD EAV sizes. You can define CKD volumes with up to 1,182,006 cylinders, about 1 TB on the DS8870.

If the number of cylinders that you specify is not an exact multiple of 1113 cylinders, then some space in the last allocated extent is wasted. For example, if you define 1114 or 3340 cylinders, 1112 cylinders are wasted. For maximum storage efficiency, consider allocating volumes that are exact multiples of 1113 cylinders. In fact, multiples of 3339 cylinders should be considered for future compatibility. If you want to use the maximum number of cylinders for a volume (that is 1,182,006 cylinders), you are not wasting cylinders, because it is an exact multiple of 1113 (1,182,006 divided by 1113 is exactly 1062). This size is also an even multiple (354) of 3339, a model 3 size.

Quick initialization

The quick initialization function initializes the data logical tracks or blocks within a specified extent range on a logical volume with the appropriate initialization pattern for the host.

Normal read-and-write access to the logical volume is allowed during the initialization process. Therefore, the extent metadata must be allocated and initialized before the quick initialization function is started. Depending on the operation, the quick initialization can be started for the entire logical volume or for an extent range on the logical volume.

The quick initialization function is started for the following operations:

- Standard logical volume creation
- Standard logical volume expansion
- Standard logical volume reinitialization
- Extent space-efficient (ESE) logical volume expansion
- ESE logical volume reinitialization
- ESE logical volume extent conversion
- Track space-efficient (TSE) or compressed TSE logical volume expansion
- TSE or compressed TSE logical volume reinitialization

Chapter 3. Data management features

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

This chapter contains information about the data management features in your storage system. Use the information in this chapter to assist you in planning, ordering licenses, and in the management of your storage system data management features.

FlashCopy SE feature

The FlashCopy SE feature allocates storage space on an as-needed basis by using space on a target volume only when it actually copies tracks from the source volume to the target volume.

Without track space-efficient (TSE) volumes, the FlashCopy function requires that all the space on a target volume be allocated and available even if no data is copied there. With space-efficient volumes, FlashCopy uses only the number of tracks that are required to write the data that is changed during the lifetime of the FlashCopy relationship, so the allocation of space is on an as-needed basis. Because it does not require a target volume that is the exact size of the source volume, the FlashCopy SE feature increases the potential for a more effective use of system storage capacity.

CAUTION:

TSE repositories greater than 50 TiB (51200 GiB) are at risk of data corruption. This can result in loss of access. Remove any TSE repository greater than 50 TiB, and allocate multiple repositories instead.

FlashCopy SE is intended for temporary copies. Unless the source data has little write activity, copy duration does not last longer than 24 hours. The best use of FlashCopy SE is when less than 20% of the source volume is updated over the life of the relationship. Also, if performance on the source or target volumes is important, standard FlashCopy is strongly recommended.

You can define the space-efficiency attribute for the target volumes during the volume creation process. A space-efficient volume can be created from any extent pool that has space-efficient storage already created in it. It is recommended, but not required, that both the source and target volumes of any FlashCopy SE relationship reside on the same cluster.

If the space-efficient source and target volumes have been created and are available, they can be selected when you create the FlashCopy relationship.

Important: Space-efficient volumes are currently supported as FlashCopy target volumes only.

After a space-efficient volume is specified as a FlashCopy target, the FlashCopy relationship becomes space-efficient. FlashCopy works the same way with a space-efficient volume as it does with a fully provisioned volume. All existing copy functions work with a space-efficient volume except for the Background Copy

function (not permitted with a space-efficient target) and the Dataset Level FlashCopy function. A miscalculation of the amount of copied data can cause the space-efficient repository to run out of space, and the FlashCopy relationship fails (that is, reads or writes to the target are prevented). You can withdraw the FlashCopy relationship to release the space.

Dynamic volume expansion

Dynamic volume expansion is the capability of the DS8000 series to increase volume capacity up to a maximum size while volumes are online to a host and not in a Copy Services relationship.

Dynamic volume expansion increases the capacity of open systems and z Systems volumes, while the volume remains connected to a host system. This capability simplifies data growth by providing volume expansion without taking volumes offline.

Some operating systems do not support a change in volume size. Therefore, a host action is required to detect the change after the volume capacity is increased.

The following volume sizes are the maximum that are supported for each storage type.

- Open systems FB volumes: 4 TiB.
- z Systems CKD volume types 3390 model 9 and custom: 65520 cylinders
- z Systems CKD volume type 3390 model 3: 3339 cylinders
- z Systems CKD volume types 3390 model A: 1,182,006 cylinders

Note: Volumes cannot be in Copy Services relationships (point-in-time copy, FlashCopy SE, Metro Mirror, Global Mirror, Metro/Global Mirror, and z/OS Global Mirror) during expansion.

Count key data and fixed block volume deletion prevention

By default, DS8000 attempts to prevent volumes that are online and in use from being deleted. The DS CLI and DS Storage Manager provides an option to force the deletion of count key data (CKD) and fixed block (FB) volumes that are in use.

For CKD volumes, *in use* means that the volumes are participating in a Copy Services relationship or are in a pathgroup. For FB volumes, *in use* means that the volumes are participating in a Copy Services relationship or there is no I/O access to the volume in the last five minutes.

If you specify the **-safe** option when you delete an FB volume, the system determines whether the volumes are assigned to non-default volume groups. If the volumes are assigned to a non-default (user-defined) volume group, the volumes are not deleted.

If you specify the **-force** option when you delete a volume, the storage system deletes volumes regardless of whether the volumes are in use.

IBM Easy Tier

Easy Tier is a DS8000 series optional feature that is provided at no cost. Its capabilities include manual volume capacity rebalance, auto performance rebalancing in both homogeneous and hybrid pools, hot spot management, rank depopulation, manual volume migration, and thin provisioning support (ESE

volumes only). Easy Tier determines the appropriate tier of storage that is based on data access requirements and then automatically and nondisruptively moves data, at the subvolume or sub-LUN level, to the appropriate tier in the storage system.

Use Easy Tier to dynamically move your data to the appropriate drive tier in your storage system with its automatic performance monitoring algorithms. You can use this feature to increase the efficiency of your flash drives and flash cards and the efficiency of all the tiers in your storage system.

You can use the features of Easy Tier between three tiers of storage within a DS8870.

Easy Tier features help you to effectively manage your system health, storage performance, and storage capacity automatically. Easy Tier uses system configuration and workload analysis with warm demotion to achieve effective overall system health. Simultaneously, data promotion and auto-rebalancing address performance while cold demotion works to address capacity.

Easy Tier Server manages distributed host caching with storage server caching and hierarchical storage management. The implementation coordinates the placement of data between the host cache, the storage server cache, and the storage tiers in the storage server, and manages consistency of the data across the set of hosts that access the data.

In automatic mode, Easy Tier data in memory persists in local storage or storage in the peer server, ensuring the Easy Tier configurations are available at failover, cold start, or Easy Tier restart.

With Easy Tier Application, you can also assign logical volumes to a specific tier. This can be useful when certain data is accessed infrequently, but needs to always be highly available.

Easy Tier Application is enhanced by two related functions:

- Easy Tier Application for IBM z Systems provides comprehensive data-placement management policy support from application to storage.
- Easy Tier application controls over workload learning and data migration provides a granular pool-level and volume-level Easy Tier control as well as volume-level tier restriction where a volume can be excluded from the Nearline tier.

The Easy Tier Heat Map Transfer utility replicates Easy Tier primary storage workload learning results to secondary storage sites, synchronizing performance characteristics across all storage systems. In the event of data recovery, storage system performance is not sacrificed.

You can also use Easy Tier in automatic mode to help with the management of your ESE thin provisioning on fixed block (FB) volumes.

An additional feature provides the capability for you to use Easy Tier in manual mode for thin provisioning. Rank depopulation is supported on ranks with ESE volumes allocated (extent space-efficient) or auxiliary volumes.

Note: Use Easy Tier in manual mode to depopulate ranks that contain TSE auxiliary volumes.

Use the capabilities of Easy Tier to support:

Three tiers

Using three tiers and efficient algorithms improves system performance and cost effectiveness.

Five types of drives are managed in up to three different tiers by Easy Tier within a managed pool. The drives within a tier must be homogeneous.

- Tier 1: flash cards and flash drives
- Tier 2: SAS (10-K or 15-K RPM) disk drives
- Tier 3: Nearline (7.2-K RPM) disk drives

If both 10-K and 15-K RPM disk drives are in the same extent pool, the disk drives are managed as a single tier. The flash cards and flash drives are managed as a single tier. In both of these cases, the rank saturation for different rank types (for example, 10K RAID-5 and 15-K RAID -5) can be different. The workload rebalancing within a single tier takes the rank saturation into consideration when attempting to achieve an equal level of saturation across the ranks within a tier.

Cold demotion

Cold data (or extents) stored on a higher-performance tier is demoted to a more appropriate tier. Easy Tier is available with two-tier disk-drive pools and three-tier pools. Sequential bandwidth is moved to the lower tier to increase the efficient use of your tiers.

Warm demotion

Active data that has larger bandwidth is demoted from either tier one (flash cards and flash drives) or tier two (Enterprise) to SAS Enterprise or Nearline SAS. Warm demotion is triggered whenever the higher tier is over its bandwidth capacity. Selected warm extents are demoted to allow the higher tier to operate at its optimal load. Warm demotes do not follow a predetermined schedule.

Manual volume or pool rebalance

Volume rebalancing relocates the smallest number of extents of a volume and restripes those extents on all available ranks of the extent pool.

Auto-rebalancing

Automatically balances the workload of the same storage tier within both the homogeneous and the hybrid pool that is based on usage to improve system performance and resource use. Use the auto-rebalancing functions of Easy Tier to manage a combination of homogeneous and hybrid pools, including relocating hot spots on ranks. With homogeneous pools, systems with only one tier can use Easy Tier technology to optimize their RAID array usage.

Rank depopulations

Allows ranks that have extents (data) allocated to them to be unassigned from an extent pool by using extent migration to move extents from the specified ranks to other ranks within the pool.

Thin provisioning

Support for the use of thin provisioning is available on ESE (FB) and standard volumes. The use of TSE volumes (FB and CKD) is not supported.

Easy Tier provides a performance monitoring capability, regardless of whether the Easy Tier license feature is activated. Easy Tier uses the monitoring process to determine what data to move and when to move it when you use automatic mode.

You can enable monitoring independently (with or without the Easy Tier license feature activated) for information about the behavior and benefits that can be expected if automatic mode were enabled.

Data from the monitoring process is included in a summary report that you can download to your Windows system. Use the IBM DS8000 Storage Tier Advisor Tool application to view the data when you point your browser to that file.

Prerequisites

The following conditions must be met to enable Easy Tier:

- The Easy Tier license feature is enabled (required for both manual and automatic mode, except when monitoring is set to All Volumes).
- For automatic mode to be active, the following conditions must be met:
 - Easy Tier automatic mode monitoring is set to either All or Auto mode.
 - For Easy Tier to manage pools, the Auto Mode Volumes must be set to either Tiered Pools or All Pools.
 - For Easy Tier Server, Easy Tier monitoring must be active. Easy Tier Server does not require Easy Tier automatic mode management.

The drive combinations that you can use with your three-tier configuration, and with the migration of your ESE volumes, are Flash, Enterprise, and Nearline.

Easy Tier: automatic mode

Use of the automatic mode of Easy Tier requires the Easy Tier license feature.

In Easy Tier, both IOPS and bandwidth algorithms determine when to migrate your data. This process can help you improve performance.

Use automatic mode to have Easy Tier relocate extents to the most appropriate storage tier in a hybrid pool, which is based on usage. Because workloads typically concentrate I/O operations on only a subset of the extents within a volume or LUN, automatic mode identifies the subset of the frequently accessed extents and relocates them to the higher-performance storage tier.

Subvolume or sub-LUN data movement is an important option to consider in volume movement because not all data at the volume or LUN level becomes hot data. For any workload, there is a distribution of data that is considered either hot or cold, which can result in significant overhead that is associated with moving entire volumes between tiers. For example, if a volume is 1 TB, you do not want to move the entire 1 TB volume when the generated heat map indicates that only 10 GB is considered hot. This capability uses your higher performance tiers to reduce the number of drives that you need to optimize performance.

Using automatic mode, you can use high performance storage tiers with a much smaller cost. This means that you invest a small portion of storage in the high-performance storage tier. You can use automatic mode for relocation and tuning without the need for your intervention, generating cost-savings while optimizing storage performance.

You also have the option of assigning specific logical volumes to a storage tier. This is useful to ensure that critical data is always highly available, regardless of how often the data is accessed.

Three-tier automatic mode is supported by the following Easy Tier functions:

- Support for ESE volumes with the thin provisioning of your FB volumes.
- Support for a matrix of device (DDM) and adapter types
- Monitoring of both bandwidth and IOPS limitations
- Data demotion between tiers
- Automatic mode hot spot rebalancing, which applies to the following auto performance rebalance situations:
 - Redistribution within a tier after a new rank is added into a managed pool
 - Redistribution within a tier after a rank is removed from a managed pool
 - Redistribution when the workload is imbalanced on the ranks within a tier of a managed pool.
- Logical volume assignment to specific storage tiers by using Easy Tier Application.
- Heat map transfer to secondary storage by using the Heat Map Transfer Utility.

To help manage and improve performance, Easy Tier is designed to identify hot data at the subvolume or sub-LUN (extent) level, which is based on ongoing performance monitoring, and then automatically relocate that data to an appropriate storage device in an extent pool that is managed by Easy Tier. Easy Tier uses an algorithm to assign heat values to each extent in a storage device. These heat values determine on what tier the data would best reside, and migration takes place automatically. Data movement is dynamic and transparent to the host server and to applications by using the data.

By default, automatic mode is enabled when the Easy Tier license feature is activated. You can temporarily disable automatic mode.

Easy Tier provides capabilities to support the automatic functions of auto-rebalance, warm demotion, and cold demotion. This includes support for pools with three tiers: Flash, Enterprise disk drives, and Nearline disk drives.

With Easy Tier you can use automatic mode to help you manage the thin provisioning of your ESE FB volumes.

Auto-rebalance

Rebalance is a function of Easy Tier automatic mode to balance the extents in the same tier that is based on usage. Auto-rebalance supports single managed pools and hybrid pools. You can use the Storage Facility Image (SFI) control to enable or disable the auto-rebalance function on all pools of an SFI. When you enable auto-rebalance, every standard and ESE volume is placed under Easy Tier management for auto-rebalancing procedures. Using auto-rebalance gives you the advantage of these automatic functions:

- Easy Tier operates within a tier, inside a managed storage pool.
- Easy Tier automatically detects performance skew and rebalances extents within the same tier.
- Easy Tier automatically rebalances extents when capacity is added to the extent pool.

In any tier, placing highly active (hot) data on the same physical rank can cause the hot rank or the associated device adapter (DA) to become a performance bottleneck. Likewise, over time skews can appear within a single tier that cannot be addressed by migrating data to a faster tier alone, and require some degree of

workload rebalancing within the same tier. Auto-rebalance addresses these issues within a tier in both hybrid and homogeneous pools. It also helps the system respond in a more timely and appropriate manner to overloading, skews, and any under-utilization that can occur from the addition or deletion of hardware, migration of extents between tiers, changes in the underlying volume configurations, and variations in the workload. Auto-rebalance adjusts the system to continuously provide optimal performance by balancing the load on the ranks and on DA pairs.

Easy Tier provides support for auto-rebalancing within homogeneous pools. If you set the Easy Tier Automatic Mode Migration control to Manage All Extent Pools, pools with a single-tier can rebalance the intra-tier ranks. If Easy Tier is turned off, then no volumes are managed. If Easy Tier is on, it manages all the volumes that it supports (Standard or ESE). TSE volumes are not supported by auto-rebalancing.

Notes:

- Standard and ESE volumes are supported.
- Merging pools are restricted to allow repository auxiliary volumes only in a single pool.
- If Easy Tier's Automatic Mode Migration control is set to Manage All Extent Pools, then single-tier pools are also managed to rebalance intra-tier ranks.

Warm demotion

Warm demotion operation demotes warm (or mostly sequential-accessed) extents in flash cards or flash drives to HDD, or from Enterprise SAS DDMs to NearLine SAS DDMs to protect the drive performance on the system. The ranks being demoted to are selected randomly. This function is triggered when bandwidth thresholds are exceeded. This means that extents are warm-demoted from one rank to another rank among tiers when extents have high bandwidth but low IOPS.

It is helpful to understand that warm demotion is different from auto-rebalancing. While both warm demotion and auto-rebalancing can be event-based, rebalancing movement takes place within the same tier while warm demotion takes place among more than one tier. Auto-rebalance can initiate when the rank configuration changes. It also periodically checks for workload that is not balanced across ranks. Warm demotion initiates when an overloaded rank is detected.

Cold demotion

Cold demotion recognizes and demotes cold or semi-cold extents to an appropriate lower-cost tier. Cold extents are demoted in a storage pool to a lower tier if that storage pool is not idle.

Cold demotion occurs when Easy Tier detects any of the following scenarios:

- Extents in a storage pool become inactive over time, while other data remains active. This is the most typical use for cold demotion, where inactive data is demoted to the SATA tier. This action frees up extents on the enterprise tier before the extents on the SATA tier become hot, helping the system be more responsive to new, hot data.
- All the extents in a storage pool become inactive simultaneously due to either a planned or unplanned outage. Disabling cold demotion assists you in scheduling extended outages or experiencing outages without effecting the extent placement.

- All extents in a storage pool are active. In addition to cold demotion by using the capacity in the lowest tier, an extent is selected which has close to zero activity, but with high sequential bandwidth and low random IOPS for the demotion. Bandwidth available on the lowest tier is also used.

All extents in a storage pool can become inactive due to a planned non-use event, such as an application that reaches its end of life. In this situation, cold demotion is disabled and you can select one of the following options:

- Allocate new volumes in the storage pool and plan on those volumes that become active. Over time, Easy Tier replaces the inactive extents on the enterprise tier with active extents on the SATA tier.
- Depopulate all of the enterprise HDD ranks. When all enterprise HDD ranks are depopulated, all extents in the pool are on the SATA HDD ranks. Store the extents on the SATA HDD ranks until they need to be deleted or archived to tape. After the enterprise HDD ranks are depopulated, move them to a storage pool.
- Leave the extents in their current locations and reactivate them later.

Figure 9 illustrates all of the migration types that are supported by the Easy Tier enhancements in a three-tier configuration. The auto-performance rebalance might also include more swap operations.

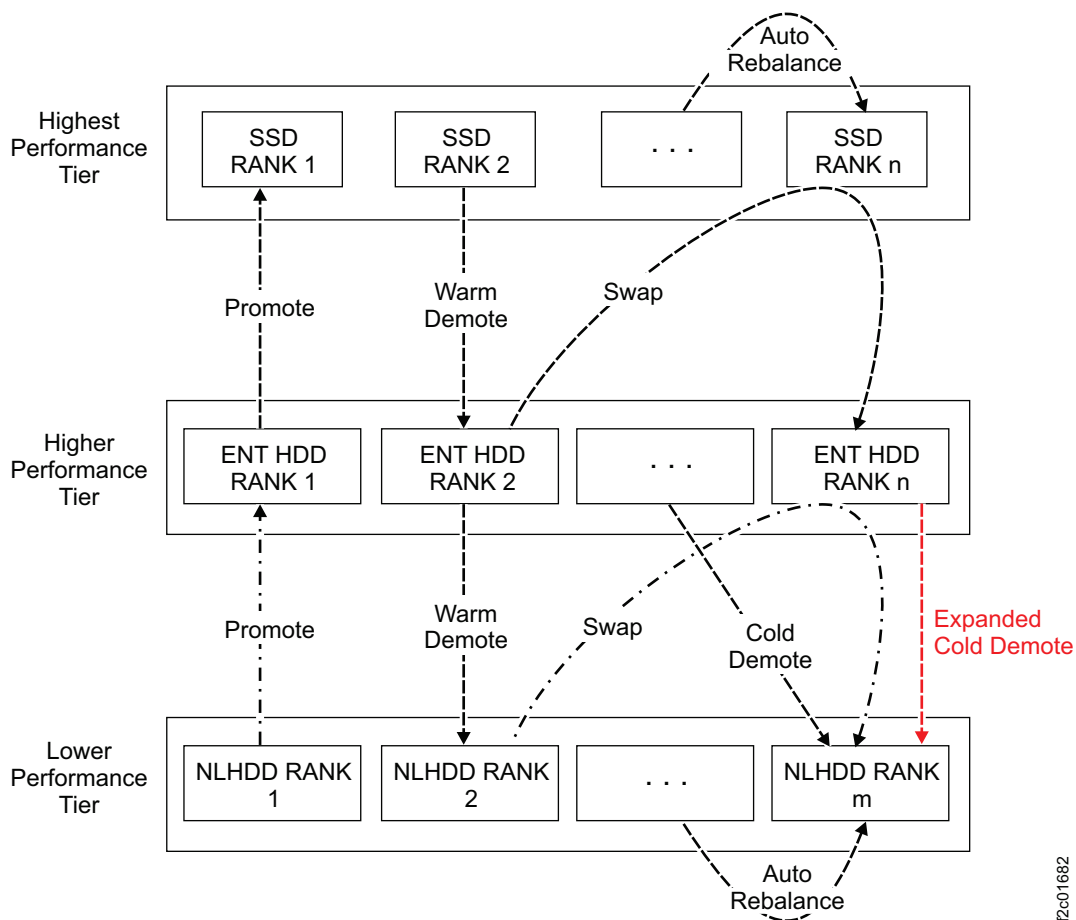


Figure 9. Three-tier migration types and their processes

Easy Tier Application

You can assign logical volumes to specific storage tiers (for non-TSE volumes). This enables applications or storage administrators to proactively influence data placement in the tiers. Applications, such as databases, can optimize access to critical data by assigning the associated logical volumes to a higher performance tier. Storage administrators, as well, can choose to assign a boot volume (for example) to a higher performance tier.

Assigning a logical volume applies to all extents that are allocated to the logical volume. Any extents added to a logical volume by dynamic extent relocation or volume expansion are also assigned to the specified tier. All assignments have an infinite lease. Assigning a volume across multiple tiers is not supported.

The completion of a logical volume assignment is a best-effort service that is based on the following Easy Tier priorities:

1. System Health

Easy Tier monitors devices to ensure that they are not overloaded for the current configuration and workload. Warm Demote operations and extra checks receive the highest priority processing in this regard.

2. Performance

Logical volume assignment requests are performed on the appropriate device types based on configuration, device capabilities, and workload characteristics.

3. Capacity

System capacity requirements are monitored.

Additionally, assignment requests can originate at multiple sources and can be delivered to Easy Tier through channels that do not guarantee ordering of messages. For this reason, the order of servicing volume assignment requests cannot be guaranteed.

Because system health is the highest priority, a logical volume assignment can be overridden by a migration operation (such as, Warm Demote), or by DS8000 microcode. As a result, although Easy Tier Application is designed to achieve eventual consistency of operations, there is no system state guarantee for an assignment, even for completed requests. The status of a logical volume assignment request can be:

- Failure

The request command is invalid and cannot be complete. A failure response is returned to the calling function.

- Transient State

The request cannot currently be completed, but is awaiting processing. A request that completed can revert to a pending state if any of its actions are undone by a higher priority request (such as a Warm Demote operation).

Additionally, the threshold (maximum capacity) for assigning logical volumes to a specified tier can be reached. The threshold is 80% of the total capacity available on that tier. In this case, all assignment requests for that tier remain pending until the assignments fall below the threshold.

- Assignment Failure

In some situations, a volume assignment request is acknowledged by Easy Tier Application, but subsequent system state changes require that the Easy Tier Application return the request as a volume assignment failure. Possible scenarios are:

- A tier definition change due to rank addition, deletion, depopulation, or merging the extent pool.
- Easy Tier automatic mode is disabled for the volume.

The assignment failure remains, until you unassign the volume. However, even while in assignment failure status, the volume is still managed by EasyTier auto functions based on its heat map.

If a logical volume is deleted, Easy Tier Application unassigns the volume, but does not identify the status as an assignment failure.

Note: For Version 7 Release 4, assignment failure works differently with the introduction of Easy Tier Application for IBM z Systems. A new status indicator, "assign pending hardware condition," is used to describe the following conditions. If a condition is later resolved, the assignment continues to be processed.

Easy Tier automatic mode becomes disabled:

The assignment remains in a pending state, and you receive a status of "assign pending hardware condition" instead of an "assign fail." If you later activate Easy Tier, the committed assignment automatically proceeds.

Target tier becomes unavailable:

You receive a status of "assign pending hardware condition," and the assignment remains in a pending state. If you later add ranks to the target tier, the committed assignment automatically proceeds.

Tier definition changes:

The physical tier is remembered, and a tier definition change does not impact the assignment.

Before Version 7 Release 4, for all the assignment failures described above, even though the condition is later resolved, the affected volumes still stay in the "assign failure" state. You need to send an unassign request to fix it. In Version 7 Release 4, you can still expect assignment failures caused by various conditions (the target tier does not exist; Easy Tier management functions are turned off; the 80% capacity limitation is exceeded; and so on) that cause the assign command to be rejected. However, after the conditions are fixed and an assign command is accepted, any changes that affect assignment activities produce only an "assign pending hardware condition," rather than an assignment-request failure.

Logical volume assignment state and request information are regularly saved to local storage or storage on the peer server. If interruptions or error conditions occur on the storage system, this data is automatically restored from the persistent storage.

Easy Tier Application for IBM z Systems

Easy Tier Application for IBM z Systems provides comprehensive data-placement management policy support between an application and storage. With this feature, you need to program the policy only once, and it is then enforced automatically. With hints about the data usage and performance expectations, storage is automatically optimized towards higher performance and efficiency. At the same time, the hint semantics relieve the application from the burden of storage resource management.

Easy Tier Application Control at pool and volume levels

Easy Tier Application Control at the pool and volume levels provides a more granular and flexible control of workload learning and data movement, as well as providing volume-level tier restriction where a volume can be excluded from the Nearline tier.

Before this feature, Easy Tier provided control at the system level. To prevent or control the placement of data, you had to disable and enable Easy Tier for the entire DS8000. Flexibility was limited. For example, if there was a performance problem within a single pool or volume, Easy Tier for the entire DS8000 needed to be stopped until the problem was corrected. This stoppage resulted in a loss of performance benefits in other pools or volumes.

Note: System-level control always has higher priority than the pool-level and volume-level control settings. If any of the system-level control settings (Easy Tier monitor; Easy Tier management) are changed, the pool and volume level control settings are reset. Changes to the system-level control settings are detected by Easy Tier every five minutes.

Several scenarios of how you can use Easy Tier customer control at the pool level and volume level are described in Table 15.

Table 15. Scenarios for Easy Tier Customer control at pool and volume levels

Function	Scenario
Suspend/resume Easy Tier learning	<p>At the pool level</p> <ul style="list-style-type: none"> A bank has a monthly and quarterly temporary batch workload, during which the workload differs from normal workloads. During the temporary batch workload, Easy Tier moves data to get good performance. However, the data configuration might not be optimal for normal workloads, so when the normal workload starts again, the performance is not as good as before. In this case, you can suspend pool learning with a duration setting when you start the temporary batch workload. After the duration expires, the pool learning resumes automatically, which makes the control easier. Alternately, you can resume the pool learning manually. You could similarly use Easy Tier control at the pool level for other tasks that have workloads that differ from normal workloads. Examples of such one-off tasks are restoring a database from a backup, database loading, and database reorganization.
	<p>At the volume level</p>
	<ul style="list-style-type: none"> One application is running the Monday-through-Friday workload, and another application is running the Saturday-through-Sunday workload. During the first workload, the application gets good performance because Easy Tier recognizes that it is hot and promotes it to SSD. But during the weekend, the first workload is no longer hot, and Easy Tier might swap another application into SSD. On the next Monday morning, the application that depends on the Monday-through-Friday workload might encounter a performance impact because Easy Tier needs time to readjust the data placement for it. In this case, you can suspend the volume learning (keep the heat) of that application at the end of the Monday-through-Friday period with an additional 48 hours of lease time. On the next Monday morning, the learning resumes automatically and the performance should be stable. During the application maintenance, such as code upgrade or backup, the performance statistics are not reflecting the real workload. To avoid the population of the learning data, in this case you can suspend the learning when doing upgrade and resume it after the upgrade is done.

Table 15. Scenarios for Easy Tier Customer control at pool and volume levels (continued)

Function	Scenario
Reset Easy Tier learning	At the pool level
	<ul style="list-style-type: none"> • You want to redefine the use of all the volumes within a storage pool. The original learning data of the volumes in the pool are no longer relevant, but you can reset the Easy Tier learning in the pool, so that Easy Tier learning reacts to the new workload quickly. Note: In many environments (especially open systems), a pool-level reset of learning is less typical as there is likely to be a mix of applications and workloads. However, this is effectively a volume-level reset of learning for all volumes in the pool. • Another scenario is when you transport a workload. You select a target pool, and the target pool’s learning data is no longer relevant. But you can reset the pool learning to react to the new workload quickly.
	At the volume level
	<ul style="list-style-type: none"> • When an application with a large amount of hot data is no longer used, the heat of the volumes associated with the application might take time to cool and stop other applications from leveraging the flash drive quickly. In this case, you can reset the learning history of the specific volumes, so other data can take advantage of the flash drive quickly. • During a database reorganization, the hot-data indexes are moved to another location by the database, so the learning history of the original data location is no longer relevant. In this case, you can reset the learning. • When you deploy a new application, you might define the file system, migrate the application data, and complete some testing before putting the new application online. The learning data during the deployment might create data-storage “noise” for the normal production workload. To alleviate the noise, you can reset the learning before the application goes online so that Easy Tier reacts quickly to the presence of the application.
Suspend/resume extent relocation	At the pool level
	<ul style="list-style-type: none"> • In one scenario, there might be some response time-sensitive period during which you want to prevent any data-movement impact to the performance. In this case, you can suspend the pool migration with a duration setting. After the duration is expired, the pool migration resumes automatically, which makes the control easier. You can also resume it manually. • In another scenario, there is a performance issue in a pool, and you want to analyze the problem. You can prevent an impact to storage during your analysis by suspending the pool’s migration to stabilize the performance of the storage.
	At the volume level
	Not applicable.
Query pool-level and volume-level Easy Tier control state	You can query the Easy Tier control state of a pool or volume.
Exclude from Nearline tier control	At the pool level
	Not applicable.
	At the volume level
	<ul style="list-style-type: none"> • If there is an application for which you do not want the data of the volume to be moved to the Nearline tier, you can exclude the volume from the Nearline tier. • During the deployment of an application, before the workload starts, the volumes that are allocated for the application might be idle. You can exclude the idle volumes from being demoted to the Nearline tier to avoid the performance issues when application starts. • To more efficiently prevent a volume from ever being demoted to the Nearline drives, you can exclude the volume from the Nearline tier so that it is only assigned to the non-Nearline tiers.

Easy Tier Server

Easy Tier Server is a unified storage caching and tiering solution across AIX servers and SAN storage.

Easy Tier Server enables the caching of storage-system data in multiple hosts based on statistics that are gathered in both the host and storage system. Easy Tier enables placing a copy of the most frequently accessed (hot) data on direct-attached storage (DAS) flash drawers on the host. Data can be read directly from flash memory that is attached to the host cache rather than from disk drives in the DS8870 storage system. This data-retrieval optimization results in improved performance, with I/O requests that are satisfied in microseconds.

The Easy Tier Server implementation consists of two major components:

- The Easy Tier Server coherency server, which runs on the DS8870 system.
The Easy Tier Server coherency server manages how data is placed on the SAN storage tiers and the SAN caches. The coherency server asynchronously communicates with the host system (the coherency client) and generates caching advice for each coherency client, based on Easy Tier placement and statistics.
- The Easy Tier Server coherency client, which runs on the host system.
The Easy Tier Server coherency client maintains local caches on DAS solid-state drives. The coherency client works independently to cache the I/O streams of its applications, providing a real-time performance enhancement. The coherency client uses the Easy Tier Server protocol to establish system-aware caching that interfaces with the coherency server.

The Easy Tier Server coherency server and the Easy Tier Server coherency client work together as follows:

- The coherency client issues SCSI commands to a logical volume to asynchronously communicate with the coherency server.
- The coherency server generates frequency-based advice that is based on a unified view of the SAN storage, hosts, and their access patterns.
- The coherency client determines what to cache, what to keep in cache, and what to evict, based on access patterns, coherency server-generated advice, and local statistics (recency- and frequency-based statistics).
- The coherency client combines the coherency server advice with its own population list, resulting in both short-term and longer-term cache population, higher hit ratio, and better storage solution optimization, including application-aware storage.

To use Easy Tier Server functions, the Easy Tier Server LIC feature must be installed and enabled on your storage system.

Easy Tier: manual mode

Easy Tier in manual mode provides the capability to migrate volumes and merge pools, under the same DS8870 system, concurrently with I/O operations.

In Easy Tier manual mode, you can dynamically relocate a logical volume between pools or within a pool to change the extent allocation method of the volume or to redistribute the volume across new ranks. This capability is referred to as *dynamic volume relocation*. You can also merge two existing pools into one without affecting the data on the logical volumes that are associated with the pools.

Enhanced functions of Easy Tier manual mode offer more capabilities. You can use manual mode to relocate your extents, or to relocate an entire volume from one pool to another pool. Later, you might also need to change your storage media or configurations. Upgrading to a new disk drive technology, rearranging the storage space, or changing storage distribution within a specific workload are typical operations that you can complete with volume relocations. Use manual mode to achieve these operations with minimal performance impact and to increase the options you have in managing your storage.

Functions and features of Easy Tier: manual mode

This section describes the functions and features of Easy Tier in manual mode.

Volume migration

Volume migration for restriping can be achieved by:

- Restriping - Relocating a subset of extents within the volume for volume migrations within the same pool.
- Rebalancing - Redistributing the volume across available ranks. This feature focuses on providing pure striping, without requiring preallocation of all the extents. This means that you can use rebalancing when only a few extents are available.

You can select which logical volumes to migrate, based on performance considerations or storage management concerns. For example, you can:

- Migrate volumes from one pool to another. You might want to migrate volumes to a different pool that has more suitable performance characteristics, such as different disk drives or RAID ranks. For example, a volume that was configured to stripe data across a single RAID can be changed to stripe data across multiple arrays for better performance. Also, as different RAID configurations become available, you might want to move a logical volume to a different pool with different characteristics, which changes the characteristics of your storage. You might also want to redistribute the available disk capacity between pools.

Notes:

- When you initiate a volume migration, ensure that all ranks are in the configuration state of Normal in the target pool.
- Volume migration is supported for standard and ESE volumes. There is no direct support to migrate auxiliary volumes. However, you can migrate extents of auxiliary volumes as part of ESE migration or rank depopulation.
- Ensure that you understand your data usage characteristics before you initiate a volume migration.
- The overhead that is associated with volume migration is comparable to a FlashCopy operation that run as a background copy.
- Change the extent allocation method that is assigned to a volume. You can relocate a volume within the same pool but with a different extent allocation method. For example, you might want to change the extent allocation method to help spread I/O activity more evenly across ranks. If you configured logical volumes in an pool with fewer ranks than now exist in the pool, you can use Easy Tier to manually redistribute the volumes across new ranks.

Note: If you specify a different extent allocation method for a volume, the new extent allocation method takes effect immediately.

Manual volume rebalance by using volume migration

Volume and pool rebalancing are designed to redistribute the extents of volumes within a non managed pool. This means skew is less likely to occur on the ranks.

Notes:

- Manual rebalancing is not allowed in hybrid or managed pools.
- Manual rebalancing is allowed in homogeneous pools.
- You cannot mix fixed block (FB) and count key data (CKD) drives.

Volume rebalance can be achieved by initiating a manual volume migration. Use volume migration to achieve manual rebalance when a rank is added to a pool, or when a large volume with rotate volumes EAM is deleted. Manual rebalance is often referred to as capacity rebalance because it balances the distribution of extents without factoring in extent usage. When a volume migration is targeted to the same pool and the target EAM is rotate extent, the volume migration acts internally as a volume rebalance.

Use volume rebalance to relocate the smallest number of extents of a volume and restripe the extents of that volume on all available ranks of the pool where it is located. The behavior of volume migration, which differs from volume rebalance, continues to operate as it did in the previous version of Easy Tier.

Notes: Use the latest enhancements to Easy Tier to:

- Migrate ESE logical volumes
- Rebalance pools by submitting a volume migration for every standard and ESE volume in a pool
- Merge pools with virtual rank auxiliary volumes in both the source and destination pool

Pools You can merge homogenous and hybrid pools. Merged pools can have 1, 2 or 3 tiers and are managed appropriately by Easy Tier in automatic mode.

Rank depopulation

Easy Tier provides an enhanced method of rank depopulation, which can be used to replace old drive technology, reconfigure pools and tear down hybrid pools. This method increases efficiency and performance when you replace or relocate whole ranks. Use the latest enhancements to Easy Tier to effect rank depopulation on any ranks in the various volume types (ESE logical, virtual rank auxiliary, TSE repository auxiliary, SE repository auxiliary, and non SE repository auxiliary).

Use rank depopulation to concurrently stop by using one or more ranks in a pool. You can use rank depopulation to do any of the following functions:

- Swap out old drive technology
- Reconfigure pools
- Tear down hybrid pools
- Change RAID types

Note: Rank depopulation is supported on ranks that have extent space efficient (ESE) extents.

Volume data monitoring

The IBM Storage Tier Advisory tool collects and reports volume data. It provides performance monitoring data even if the license feature is not activated.

You can monitor the use of storage at the volume extent level using the monitoring function. . Monitoring statistics are gathered and analyzed every 24 hours. In an Easy Tier managed pool, the analysis is used to form an extent relocation plan for the pool, which provides a recommendation, which is based on your current plan, for relocating extents on a volume to the most appropriate storage device. The results of this data is summarized in a report that you can download. For more information, see “Storage Tier Advisor tool” on page 68.

Table 16 describes monitor settings and mirrors the monitor settings in the DS CLI.

Table 16. Monitoring settings for the Easy Tier license feature

Monitor Setting	Not installed	Installed
All Volumes	All volumes are monitored.	All volumes are monitored.
Auto Mode Volumes	No volumes are monitored.	Volumes in pools that are managed by Easy Tier are monitored.
No Volumes	No volumes are monitored.	No volumes are monitored.

The default monitoring setting for Easy Tier **Auto Mode** is On. Volumes in managed pools are monitored when the Easy Tier license feature is activated. Volumes are not monitored if the Easy Tier license feature is not activated.

You can determine whether volumes are monitored and also disable the monitoring process temporarily, by using either the DS CLI or the DS8000 Storage Management GUI.

Easy Tier Heat Map Transfer Utility

A *heat map* is a workload activity metric that is calculated for each extent in a logical volume. The workload activity is expressed as a temperature gradient from hot (high activity) to cold (low activity). Use of the heat map transfer utility requires the Easy Tier monitoring function to be enabled at each of the primary and secondary storage systems that are involved in the heat map transfer.

The heat map transfer utility periodically transfers Easy Tier heat map information from primary to secondary storage systems. The secondary storage system generates migration plans based on the heat map data and (the secondary storage system's) current physical configuration. In this way, the performance characteristics of the secondary storage are consistently updated to reflect that of primary storage. Multiple secondary storage systems are supported. Alternatively, you can have multiple primary storage systems that are associated with a single secondary storage system. It is recommended that the secondary storage system has the same physical configuration as the primary storage system. Secondary storage systems are then workload optimized based on primary storage system usage, with no performance penalties if data recovery is necessary.

Note: Currently, the heat map transfer utility does not support replicating tier assignment instructions of the Easy Tier Application from the primary to secondary storage systems. To reflect the same tier assignment on the secondary storage systems, issue the same tier assignment commands on the secondary storage systems.

Data that occurs in the I/O cache layer (including the storage and server-side cache) is not monitored by Easy Tier and not reflected in an Easy Tier heat map.

If a workload failover occurs, the secondary storage system:

- Uses the heat map data that is transferred from the primary storage system.
- Maintains performance levels equivalent to the primary storage system while the primary storage system is unavailable.

Note: Without the same physical configuration, a secondary storage site is able to replicate the heat map data, but is unlikely to be able to replicate the performance characteristics of the primary storage system.

The heat map transfer utility runs either on a separate Windows or Linux host, or on Tivoli Storage Productivity Center for Replication. From the host, the heat map transfer utility accesses the primary and secondary storage sites by using an out-of-band IP connection. Transfer of heat map data occurs through the heat map transfer utility host, as illustrated in Figure 10.

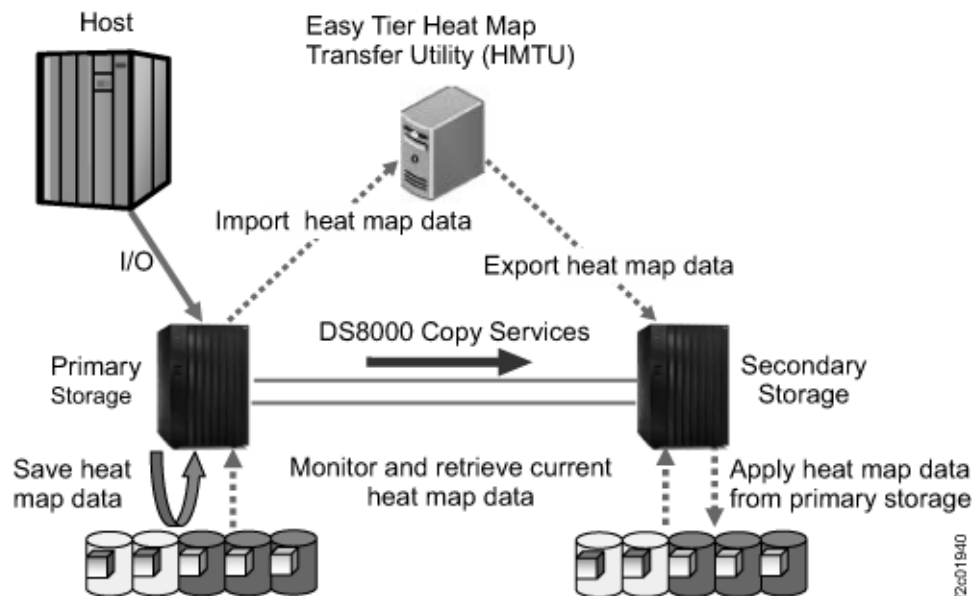


Figure 10. Flow of heat map data

The heat map transfer utility imports the heat map data from the primary storage system, and analyzes this data to:

- Identify those volumes that have a peer-to-peer remote client (PPRC) relationship.
- Determine the type of PPRC relationship that exists. The relationship can be Metro Mirror, Global Copy, Global Mirror, or Metro Global Mirror.

In a Metro Global Mirror environment, DS8000 storage systems can be added under the heat map transfer utility management. Under this management, the heat map transfer utility treats the systems as Metro Mirror plus Global Mirror (Global Copy and FlashCopy) relationships. The utility detects the Metro Mirror and Global Mirror relationships automatically and performs the heat map data transfer for the relationships on the systems separately.

There are restrictions in a heat map transfer in Metro Global Mirror environment. For example, assume volumes A, B, C and D, where:

- Volume A is the Metro Mirror primary (or source) volume
- Volume B is the Metro Mirror secondary (or target) volume and Global Mirror primary volume at the same time.
- Volume C is the Global Mirror secondary volume and FlashCopy source volume at the same time. The FlashCopy target volume is referred to as the D volume.
 - Heat map data is transferred only from volumes A and B and volumes B and C. No heat map data is transferred to the volume D copy or any additional test copies that you create.
 - Heat map data that is transferred to volume C might lag for a maximum of 36 hours from volume A. After the transfer to volumes A and B is complete, it might take a maximum of 24 hours (the default Easy Tier heat map data generation interval) for volume B to generate heat map data. There is a 12-hour interval (the default heat map transfer interval) for the volumes B and C data transfer.

The heat map information for the selected volumes is then periodically copied from the primary storage system to the heat map transfer utility host (default copy period is 12 hours). The heat-map-transfer utility determines the target secondary storage system that is based on PPRC volume mapping. The utility transfers the heat-map data to the associated secondary storage systems. The heat-map data is then imported to the secondary storage system, and Easy Tier migration plans are generated based on the imported and existing heat map. Finally, the result of the heat map transfer is recorded (in memory and to a file).

To enable heat map transfer, the heat-map transfer-control switch that is on the secondary storage system needs to be enabled `-ethmtmode enabled`. This is the default mode. Use the DSCLI command `chsi` to enable or disable heat map transfer:

```
chsi -ethmtmode enable | disable
```

The scope of heat map transfer is determined by the Easy Tier automatic mode setting:

- To automatically transfer the heat map data and manage data placement for logical volumes in multi-tiered pools, use the Easy Tier control default settings (`-etmonitor automode, -etautomode tiered`).
- To automatically transfer the heat map data and manage data placement for logical volumes in all pools, use the Easy Tier control settings (`-etmonitor all, -etautomode all`).

Note: For PPRC relationships by using Global Mirror, Easy Tier manages data placement of the Global Copy target and FlashCopy source only, and does not manage data placement for a FlashCopy target that is involved in the Global Mirror relationship.

If you do not have an Easy Tier license, and want to run an Easy Tier evaluation on both the primary and secondary storage systems, set the Easy Tier control on both storage systems to "monitor only" (`-etmonitor all`). The heat map transfer utility then automatically transfers the heat map data and uses this data to generate an Easy Tier report, without changing the data layout on either of the storage systems.

Migration process management

You can initiate volume migrations and pause, resume, or cancel a migration process that is in progress.

Volumes that are eligible for migration are dependent on the state and access of the volumes. Table 17 shows the states that are required to allow migration with Easy Tier.

Table 17. Volume states required for migration with Easy Tier

	Volume state	Is migration allowed with Easy Tier?
Access state	Online	Yes
	Fenced	No
Data state	Normal	Yes
	Pinned	No
	Read only	Yes
	Inaccessible	No
	Indeterminate data loss	No
	Extent fault	No

Initiating volume migration

With Easy Tier, you can migrate volumes from one extent pool to another. The time to complete the migration process might vary, depending on what I/O operations are occurring on your storage unit.

If an error is detected during the migration process, the storage facility image (SFI) attempts the extent migration again after a short time. If an extent cannot be successfully migrated, the migration is stopped, and the configuration state of the logical volume is set to migration error.

Pausing and resuming migration

You can pause volumes that are being migrated. You can also resume the migration process on the volumes that were paused.

Canceling migration

You can cancel the migration of logical volumes that are being migrated. The volume migration process pre-allocates all extents for the logical volume when you initiate a volume migration. All pre-allocated extents on the logical volume that are not migrated are released when you cancel a volume migration. The state of the logical volumes changes to migration-canceled and the target extent pool that you specify on a subsequent volume migration is limited to either the source extent pool or target extent pool of the original volume migration.

Note: If you initiate a volume migration but the migration was queued and not in progress, then the cancel process returns the volume to normal state and not migration-canceled.

Storage Tier Advisor tool

IBM DS8000 Storage Tier Advisor Tool adds performance reporting capability to your storage system.

The Storage Tier Advisor tool is a Windows application that provides a graphical representation of performance data that is collected by Easy Tier over a 24-hour operational cycle. You can use the application to view the data when you point your browser to the file. The Storage Tier Advisor tool supports the enhancements that are provided with Easy Tier, including support for flash cards, flash drives (SSDs), Enterprise, and Nearline disk drives for DS8870 and the auto performance rebalance feature. You can download the Storage Tier Advisor Tool (ftp://software.ibm.com/storage/ds8000/updates/DS8K_Customer_Download_Files/Storage_Tier_Advisor_Tool/).

To extract the performance summary data that is generated by the Storage Tier Advisor tool, you can use the DS CLI. When you extract summary data, two files are provided, one for each server in the storage facility image (SFI server). The download operation initiates a long running task to collect performance data from both selected storage facility images. This information can be provided to IBM if performance analysis or problem determination is required.

You can view information to analyze workload statistics and evaluate which logical volumes might be candidates for Easy Tier management. If the Easy Tier feature is not installed and enabled, you can use the performance statistics that are gathered by the monitoring process to help you determine whether to use Easy Tier to enable potential performance improvements in your storage environment.

Easy Tier reporting improvements

The reporting mechanism of Easy Tier and the Storage Tier Advisor Tool that uses Easy Tier includes updates to a workload categorization, workload skew curve, and the data-movement daily report.

The output of the Storage Tier Advisor Tool (STAT) is based on data collected by the Easy Tier monitoring function. Active data moves to a flash drive (SSD) storage tier while inactive data is demoted to a nearline storage tier. Active large data is sequential I/O, which might not be suitable to a flash drive tier, while low-active data might not be active enough to be placed on a Flash tier. The reporting improvements help you analyze this type of data activity and evaluate workload statistics across the storage tiers.

The STAT utility analyzes data that Easy Tier gathers and creates a set of comma-separated value (.csv) files for the workload categorization, workload skew curve, and data-movement daily report that you can download and generate a graphical display of the data from the *.csv files. This information provides insights into your storage workload.

For information on the workload categorization, workload skew curve, and the daily data movement report, see the Easy Tier section under Product Overview in the IBM DS8000 series online product documentation (www.ibm.com/support/knowledgecenter/ST8NCA/product_welcome/ds8000_kcwelcome.html).

Easy Tier considerations and limitations

When you plan for volume migration, it is important to consider how Easy Tier functions with storage configurations, and recognize its limitations.

Migration considerations

The following information might be helpful in using Easy Tier with your DS8000 storage system:

- You cannot initiate a volume migration on a volume that is being migrated. The first migration must complete first.
- You cannot initiate, pause, resume, or cancel migration on selected volumes that are aliases or virtual volumes.
- You cannot migrate volumes from one extent pool to another or change the extent allocation method unless the Easy Tier feature is installed on the storage system.
- Volume migration is supported for standard, auxiliary, and ESE volumes.
- If you specify a different extent allocation method for a volume, the new extent allocation method takes effect immediately.
- A volume that is being migrated cannot be expanded and a volume that is being expanded cannot be migrated.
- When a volume is migrated out of an extent pool that is managed with Easy Tier, or when Easy Tier is no longer installed, the DS8870 disables Easy Tier and no longer automatically relocates high activity I/O data on that volume between storage devices.

Limitations

The following limitations apply to the use of Easy Tier:

- TSE logical volumes do not support extent migration. This limitation means that these entities do not support Easy Tier manual mode or Easy Tier automatic mode.
- You cannot merge two extent pools:
 - If both extent pools contain TSE volumes.
 - If there are TSE volumes on the flash ranks.
 - If you selected an extent pool that contains volumes that are being migrated.
- It might be helpful to know that some basic characteristics of Easy Tier might limit the applicability for your generalized workloads. The granularity of the extent that can be relocated within the hierarchy is large (1 GB). Additionally, the time period over which the monitoring is analyzed is continuous, and long (24 hours). Therefore, some workloads might have hot spots, but when considered over the range of the relocation size, they do not appear, on average, to be hot. Also, some workloads might have hot spots for short periods of time, but when considered over the duration of the analysis window, the hot spots do not appear, on average, to be hot.

VMware vStorage API for Array Integration support

The DS8870 provides support for the VMware vStorage API for Array Integration (VAAI).

The VAAI API offloads storage processing functions from the server to the DS8870, reducing the workload on the host server hardware for improved performance on both the network and host servers.

The DS8870 supports the following operations:

Atomic test and set or VMware hardware-assisted locking

The hardware-assisted locking feature uses the VMware Compare and

Write command for reading and writing the volume's metadata within a single operation. With the Compare and Write command, the DS8870 provides a faster mechanism that is displayed to the volume as an atomic action that does not require locking the entire volume.

The Compare and Write command is supported on all open systems fixed block volumes, including Metro Mirror and Global Mirror primary volumes and FlashCopy source and target volumes.

XCOPY or Full Copy

The XCOPY (or extended copy) command copies multiple files from one directory to another or across a network.

Full Copy copies data from one storage array to another without writing to the VMware ESX Server (VMware vStorage API).

The following restrictions apply to XCOPY:

- XCOPY is not supported on Extent Space Efficient (ESE) or Track Space Efficient (TSE) volumes
- XCOPY is not supported on volumes greater than 2 TB
- The target of an XCOPY cannot be a Metro Mirror or Global Mirror primary volume
- The Fixed Block FlashCopy license is required

Block Zero (Write Same)

The SCSI Write Same command is supported on all volumes. This command efficiently writes each block, faster than standard SCSI write commands, and is optimized for network bandwidth usage.

IBM vCenter plug-in for ESX 4.x

The IBM vCenter plug-in for ESX 4.x provides support for the VAAI interfaces on ESX 4.x.

For information on how to attach a VMware ESX Server host to a DS8870 with Fibre Channel adapters, visit IBM DS8000 Information Center IBM DS8000 series online product documentation (www.ibm.com/support/knowledgecenter/ST8NCA/product_welcome/ds8000_kcwelcome.html) and select **Configuring > Attaching Host > VMware ESX Server host attachment**.

VMware vCenter Site Recovery Manager 5.0

VMware vCenter Site Recovery Manager (SRM) provides methods to simplify and automate disaster recovery processes. IBM Site Replication Adapter (SRA) communicates between SRM and the storage replication interface. SRA support for SRM 5.0 includes the new features for planned migration, reprotection, and failback. The supported Copy Services are Metro Mirror, Global Mirror, Metro-Global Mirror, and FlashCopy.

The IBM Storage Management Console plug-in enables VMware administrators to manage their systems from within the VMware management environment. This plug-in provides an integrated view of IBM storage to VMware virtualize datastores that are required by VMware administrators. For information, see the IBM Storage Management Console for VMware vCenter (www.ibm.com/support/knowledgecenter/HW213_7.4.0/hsg/hsg_vcplugin_kcwelcome.html) online documentation.

Performance for IBM z Systems

The DS8000 series supports the following IBM performance enhancements for IBM z Systems environments.

- Parallel access volumes (PAVs)
- Multiple allegiance
- z/OS Distributed Data Backup
- z/HPF extended distance capability

Parallel access volumes

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

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

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

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

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

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

Multiple allegiance

With multiple allegiance, the storage unit can run concurrent, multiple requests from multiple hosts.

Traditionally, IBM storage subsystems allow only one channel program to be active to a disk volume at a time. This means that, after the subsystem accepts an I/O request for a particular unit address, this unit address appears "busy" to

subsequent I/O requests. This single allegiance capability ensures that additional requesting channel programs cannot alter data that is already being accessed.

By contrast, the storage unit is capable of multiple allegiance (or the concurrent execution of multiple requests from multiple hosts). That is, the storage unit can queue and concurrently run multiple requests for the same unit address, if no extent conflict occurs. A conflict refers to either the inclusion of a Reserve request by a channel program or a Write request to an extent that is in use.

z/OS Distributed Data Backup

z/OS Distributed Data Backup (zDDB) is an optional licensed feature that allows hosts, which are attached through a FICON or ESCON interface, to access data on fixed block (FB) volumes through a device address on FICON or ESCON interfaces.

If the zDDB LIC feature key is installed and enabled and a volume group type specifies either FICON or ESCON interfaces, this volume group has implicit access to all FB logical volumes that are configured in addition to all CKD volumes specified in the volume group. In addition, this optional feature enables data backup of open systems from distributed server platforms through a z Systems host. The feature helps you manage multiple data protection environments and consolidate those into one environment that is managed by IBM z Systems. For more information, see “z/OS Distributed Data Backup” on page 134.

z/HPF extended distance

z/HPF extended distance reduces the impact that is associated with supported commands on current adapter hardware, improving FICON throughput on the DS8000 I/O ports. The DS8000 also supports the new zHPF I/O commands for multitrack I/O operations.

Copy Services

Copy Services functions can help you implement storage solutions to keep your business running 24 hours a day, 7 days a week. Copy Services include a set of disaster recovery, data migration, and data duplication functions.

The storage system supports Copy Service functions that contribute to the protection of your data. These functions are also supported on the IBM TotalStorage Enterprise Storage Server®.

Notes:

- If you are creating paths between an older release of the DS8000 (Release 5.1 or earlier), which supports only 4-port host adapters, and a newer release of the DS8000 (Release 6.0 or later), which supports 8-port host adapters, the paths connect only to the lower four ports on the newer storage system.
- The maximum number of FlashCopy relationships that are allowed on a volume is 65534. If that number is exceeded, the FlashCopy operation fails.
- The size limit for volumes or extents in a Copy Service relationship is 2 TB.
- Thin provisioning functions in open-system environments are supported for the following Copy Services functions:
 - FlashCopy relationships

- Global Mirror relationships if the Global Copy A and B volumes are Extent Space Efficient (ESE) volumes. The FlashCopy target volume (Volume C) in the Global Mirror relationship can be an ESE volume, Target Space Efficient (TSE) volume, or standard volume.
- PPRC supports any intermix of T10-protected or standard volumes. FlashCopy does not support intermix.

The following Copy Services functions are available as optional features:

- Point-in-time copy, which includes IBM FlashCopy and Space-Efficient FlashCopy

The FlashCopy function enables you to make point-in-time, full volume copies of data, so that the copies are immediately available for read or write access. In z Systems environments, you can also use the FlashCopy function to perform data set level copies of your data.

- Remote mirror and copy, which includes the following functions:

- Metro Mirror

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

- Global Copy

Global Copy is a nonsynchronous long-distance copy function where incremental updates are sent from the local to the remote site on a periodic basis.

- Global Mirror

Global Mirror is a long-distance remote copy function across two sites by using asynchronous technology. Global Mirror processing is designed to provide support for unlimited distance between the local and remote sites, with the distance typically limited only by the capabilities of the network and the channel extension technology.

- Metro/Global Mirror (a combination of Metro Mirror and Global Mirror)

Metro/Global Mirror is a three-site remote copy solution. It uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

- Multiple Target PPRC

Multiple Target PPRC builds and extends the capabilities of Metro Mirror and Global Mirror. It allows data to be mirrored from a single primary site to two secondary sites simultaneously. You can define any of the sites as the primary site and then run Metro Mirror replication from the primary site to either of the other sites individually or both sites simultaneously.

- Remote mirror and copy for z Systems environments, which includes z/OS Global Mirror

Note: When FlashCopy is used on FB (open) volumes, the source and the target volumes must have the same protection type of either T10 DIF or standard.

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

You can manage these functions through a command-line interface that is called the DS CLI and a Web-based interface that is called the DS8000 Storage Management GUI. You can use the DS8000 Storage Management GUI to set up and manage the following types of data-copy functions from any point where network access is available:

Point-in-time copy (FlashCopy)

You can use the FlashCopy function to make point-in-time, full volume copies of data, with the copies immediately available for read or write access. In z Systems environments, you can also use the FlashCopy function to perform data set level copies of your data. You can use the copy with standard backup tools that are available in your environment to create backup copies on tape.

FlashCopy is an optional function. To use it, you must purchase one of the point-in-time 242x indicator feature and 239x function authorization features.

The FlashCopy function creates a copy of a source volume on the target volume. This copy is called a point-in-time copy. When you initiate a FlashCopy operation, a FlashCopy relationship is created between a source volume and target volume. A FlashCopy relationship is a *mapping* of the FlashCopy source volume and a FlashCopy target volume. This mapping allows a point-in-time copy of that source volume to be copied to the associated target volume. The FlashCopy relationship exists between the volume pair in either case:

- From the time that you initiate a FlashCopy operation until the storage system copies all data from the source volume to the target volume.
- Until you explicitly delete the FlashCopy relationship if it was created as a persistent FlashCopy relationship.

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

- The amount of data to be copied
- The number of background copy processes that are occurring
- The other activities that are occurring on the storage systems

The FlashCopy function supports the following copy options:

Consistency groups

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

Change recording

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

Establish FlashCopy on existing Metro Mirror source

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 Incremental FlashCopy

Allows a source volume to establish incremental flash copies to a maximum of 12 targets.

Multiple Relationship FlashCopy

Allows a source volume to have multiple (up to 12) target volumes at the same time.

Persistent FlashCopy

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

Refresh target volume

Refresh a FlashCopy relationship, without recopying all tracks from the source volume to the target volume.

Resynchronizing FlashCopy volume pairs

Update an initial point-in-time copy of a source volume without having to recopy your entire volume.

Reverse restore

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

Reset SCSI reservation on target volume

If there is a SCSI reservation on the target volume, the reservation is released when the FlashCopy relationship is established. If this option is not specified and a SCSI reservation exists on the target volume, the FlashCopy operation fails.

Remote Pair FlashCopy

Figure 11 on page 76 illustrates how Remote Pair FlashCopy works. If Remote Pair FlashCopy is used to copy data from Local A to Local B, an equivalent operation is also performed from Remote A to Remote B. FlashCopy can be performed as described for a Full Volume FlashCopy, Incremental FlashCopy, and Dataset Level FlashCopy.

The Remote Pair FlashCopy function prevents the Metro Mirror relationship from changing states and the resulting momentary period where Remote A is out of synchronization with Remote B. This feature provides a solution for data replication, data migration, remote copy, and disaster recovery tasks.

Without Remote Pair FlashCopy, when you established a FlashCopy relationship from Local A to Local B, by using a Metro Mirror primary volume as the target of that FlashCopy relationship, the corresponding Metro Mirror volume pair went from “full duplex” state to “duplex pending” state as long as the FlashCopy data was being transferred to the Local B. The time that it took to complete the copy of the FlashCopy data

until all Metro Mirror volumes were synchronous again, depended on the amount of data transferred. During this time, the Local B would be inconsistent if a disaster were to have occurred.

Note: Previously, if you created a FlashCopy relationship with the **Preserve Mirror, Required** option, by using a Metro Mirror primary volume as the target of that FlashCopy relationship, and if the status of the Metro Mirror volume pair was not in a “full duplex” state, the FlashCopy relationship failed. That restriction is now removed. The Remote Pair FlashCopy relationship completes successfully with the “Preserve Mirror, Required” option, even if the status of the Metro Mirror volume pair is either in a suspended or duplex pending state.




Figure 11. Remote Pair FlashCopy

Note: The DS8870 supports Incremental FlashCopy and Metro Global Mirror Incremental Resync on the same volume.

Remote mirror and copy

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

Remote mirror and copy is an optional feature that provides data backup and disaster recovery. To use it, you must purchase at least one of the remote mirror and copy 242x indicator feature and 239x function authorization features.

The remote mirror and copy feature provides synchronous (Metro Mirror) and asynchronous (Global Copy) data mirroring. The main difference is that the Global Copy feature can operate at long distances, even continental distances, with

minimal impact on applications. Distance is limited only by the network and channel extenders technology capabilities. The maximum supported distance for Metro Mirror is 300 km.

With Metro Mirror, application write performance depends on the available bandwidth. Global Copy enables better use of available bandwidth capacity to enable you to include more of your data to be protected.

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

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

The following sections describe the remote mirror and copy functions.

Synchronous mirroring (Metro Mirror)

Provides real-time mirroring of logical volumes (a source and a target) between two storage systems that can be located up to 300 km from each other. With Metro Mirror copying, the source and target volumes can be on the same storage system or on separate storage systems. You can locate the storage system at another site, some distance away.

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

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

Asynchronous mirroring (Global Copy)

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

The updates are tracked and periodically copied to the target volumes. As a consequence, there is no guarantee that data is transferred in the same sequence that was applied to the source volume.

To get a consistent copy of your data at your remote site, periodically switch from Global Copy to Metro Mirror mode, then either stop the application I/O or freeze data to the source volumes by using a manual process with freeze and run commands. The freeze and run functions can

be used with external automation software such as Geographically Dispersed Parallel Sysplex™ (GDPS®), which is available for z Systems environments, to ensure data consistency to multiple Metro Mirror volume pairs in a specified logical subsystem.

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

Suspend and resume

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

Copy out-of-synchronous data

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

Copy an entire volume or not copy the volume

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

Global Mirror

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

Global Mirror integrates both the Global Copy and FlashCopy functions.

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

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

Global Mirror supports up to 32 Global Mirror sessions per storage facility image. Previously, only one session was supported per storage facility image.

You can use multiple Global Mirror sessions to fail over only data that is assigned to one host or application instead of forcing you to fail over all data if one host or application fails. This process provides increased flexibility to control the scope of a failover operation and to assign different options and attributes to each session.

The DS CLI and DS Storage Manager display information about the sessions, including the copy state of the sessions.

Practice copying and consistency groups

To get a consistent copy of your data, you can pause Global Mirror on a consistency group boundary. Use the pause command with the secondary storage option. (For more information, see the DS CLI Commands reference.) After verifying that Global Mirror is paused on a consistency boundary (state is Paused with Consistency), the secondary storage system and the FlashCopy target storage system or device are consistent. You can then issue either a FlashCopy or Global Copy command to make a practice copy on another storage system or device. You can immediately resume Global Mirror, without the need to wait for the practice copy operation to finish. Global Mirror then starts forming consistency groups again. The entire pause and resume operation generally takes just a few seconds.

Metro/Global Mirror

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

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

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

Multiple Target PPRC

Provides an enhancement to disaster recovery solutions by allowing data to be mirrored from a single primary site to two secondary sites simultaneously. The function builds on and extends Metro Mirror and Global Mirror capabilities. Various interfaces and operating systems support the function. Disaster recovery scenarios depend on support from controlling software such as Geographically Dispersed Parallel Sysplex (GDPS) and Tivoli Storage Productivity Center for Replication.

Additional information is provided in Knowledge Center (IBM Knowledge Center website). Use the search or filtering functions or find it in the navigation by clicking **System Storage > Disk systems> Enterprise Storage Servers> DS8000**.

z/OS Global Mirror

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

The z/OS Global Mirror function is an optional function. To use it, you must purchase the remote mirror for z/OS 242x indicator feature and 239x function authorization feature.

z/OS Metro/Global Mirror Incremental Resync

z/OS Metro/Global Mirror Incremental Resync is an enhancement for z/OS Metro/Global Mirror. z/OS Metro/Global Mirror Incremental Resync can eliminate the need for a full copy after a HyperSwap® situation in 3-site z/OS Metro/Global Mirror configurations. The storage system supports z/OS Metro/Global Mirror that is a 3-site mirroring solution that uses IBM System Storage Metro Mirror and z/OS Global Mirror (XRC). The z/OS Metro/Global Mirror Incremental Resync capability is intended to enhance this solution by enabling resynchronization of data between sites by using only the changed data from the Metro Mirror target to the z/OS Global Mirror target after a HyperSwap operation.

If an unplanned failover occurs, you can use the z/OS Soft Fence function to prevent any system from accessing data from an old primary PPRC site. For more information, see the *GDPS/PPRC Installation and Customization Guide*, or the *GDPS/PPRC HyperSwap Manager Installation and Customization Guide*.

z/OS Global Mirror Multiple Reader (enhanced readers)

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

Interoperability with existing and previous generations of the DS8000 series

All of the remote mirroring solutions that are documented in the sections above use Fibre Channel as the communications link between the primary and secondary storage systems. The Fibre Channel ports that are used for remote mirror and copy can be configured as either a dedicated remote mirror link or as a shared port between remote mirroring and Fibre Channel Protocol (FCP) data traffic.

The remote mirror and copy solutions are optional capabilities and are compatible with previous generations of DS8000 series. They are available as follows:

- Metro Mirror indicator feature numbers 75xx and 0744 and corresponding DS8000 series function authorization (2396-LFA MM feature numbers 75xx)
- Global Mirror indicator feature numbers 75xx and 0746 and corresponding DS8000 series function authorization (2396-LFA GM feature numbers 75xx).

The DS8000 series systems can also participate in Global Copy solutions with the IBM TotalStorage ESS Model 750, IBM TotalStorage ESS Model 800, and IBM System Storage DS6000™ series systems for data migration. For more information on data migration and migration services, contact IBM or a Business Partner representative.

Global Copy is a non-synchronous long-distance copy option for data migration and backup, and is available under Metro Mirror and Global Mirror licenses or Remote Mirror and Copy license on older DS8000, ESS, or DS6000 systems.

Disaster recovery through Copy Services

Through Copy Services functions, you can prepare for a disaster by backing up, copying, and mirroring your data at local (production) and remote sites.

Having a disaster recovery plan can ensure that critical data is recoverable at the time of a disaster. Because most disasters are unplanned, your disaster recovery plan must provide a way to recover your applications quickly, and more importantly, to access your data. Consistent data to the same point-in-time across all storage units is vital before you can recover your data at a backup (normally your remote) site.

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

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

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

Data consistency can be achieved through the following methods:

Manually using external software (without Global Mirror)

You can use Metro Mirror, Global Copy, and FlashCopy functions to create a consistent and restartable copy at your recovery site. These functions require a manual and periodic suspend operation at the local site. For instance, you can enter the **freeze** and **run** commands with external automated software. Then, you can initiate a FlashCopy function to make a

consistent copy of the target volume for backup or recovery purposes. Automation software is not provided with the storage system; it must be supplied by the user.

Note: The **freeze** operation occurs at the same point-in-time across all links and all storage systems.

Automatically (with Global Mirror and FlashCopy)

You can automatically create a consistent and restartable copy at your intermediate or remote site with minimal or no interruption of applications. This automated process is available for two-site Global Mirror or three-site Metro / Global Mirror configurations. Global Mirror operations automate the process of continually forming consistency groups. It combines Global Copy and FlashCopy operations to provide consistent data at the remote site. A master storage unit (along with subordinate storage units) internally manages data consistency through consistency groups within a Global Mirror configuration. Consistency groups can be created many times per hour to increase the currency of data that is captured in the consistency groups at the remote site.

Note: A consistency group is a collection of session-grouped volumes across multiple storage systems. Consistency groups are managed together in a session during the creation of consistent copies of data. The formation of these consistency groups is coordinated by the master storage unit, which sends commands over remote mirror and copy links to its subordinate storage units.

If a disaster occurs at a local site with a two or three-site configuration, you can continue production on the remote (or intermediate) site. The consistent point-in-time data from the remote site consistency group enables recovery at the local site when it becomes operational.

Resource groups for Copy Services scope limiting

Resource groups are used to define a collection of resources and associate a set of policies relative to how the resources are configured and managed. You can define a network user account so that it has authority to manage a specific set of resources groups.

Copy Services scope limiting overview

Copy services scope limiting is the ability to specify policy-based limitations on Copy Services requests. With the combination of policy-based limitations and other inherent volume-addressing limitations, you can control which volumes can be in a Copy Services relationship, which network users or host LPARs issue Copy Services requests on which resources, and other Copy Services operations.

Use these capabilities to separate and protect volumes in a Copy Services relationship from each other. This can assist you with multitenancy support by assigning specific resources to specific tenants, limiting Copy Services relationships so that they exist only between resources within each tenant's scope of resources, and limiting a tenant's Copy Services operators to an "operator only" role.

When managing a single-tenant installation, the partitioning capability of resource groups can be used to isolate various subsets of an environment as if they were separate tenants. For example, to separate mainframes from distributed system servers, Windows from UNIX, or accounting departments from telemarketing.

Using resource groups to limit Copy Service operations

Figure 12 illustrates one possible implementation of an exemplary environment that uses resource groups to limit Copy Services operations. Two tenants (Client A and Client B) are illustrated that are concurrently operating on shared hosts and storage systems.

Each tenant has its own assigned LPARs on these hosts and its own assigned volumes on the storage systems. For example, a user cannot copy a Client A volume to a Client B volume.

Resource groups are configured to ensure that one tenant cannot cause any Copy Services relationships to be initiated between its volumes and the volumes of another tenant. These controls must be set by an administrator as part of the configuration of the user accounts or access-settings for the storage system.




Figure 12. Implementation of multiple-client volume administration

Resource groups functions provide additional policy-based limitations to users or the DS8000 storage systems, which in conjunction with the inherent volume addressing limitations support secure partitioning of Copy Services resources between user-defined partitions. The process of specifying the appropriate limitations is completed by an administrator using resource groups functions.

Note: User and administrator roles for resource groups are the same user and administrator roles used for accessing your DS8000 storage system. For example, those roles include storage administrator, Copy Services operator, and physical operator.

The process of planning and designing the use of resource groups for Copy Services scope limiting can be complex. For more information on the rules and policies that must be considered in implementing resource groups, see topics about

resource groups. For specific DS CLI commands used to implement resource groups, see the *IBM DS8000 Command-Line Interface User's Guide*.

Comparison of licensed functions

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

Table 18 provides a brief summary of the characteristics of the Copy Services features that are available for the storage system.

Table 18. Comparison of licensed functions

Licensed function	Description	Advantages	Considerations
Multiple Target PPRC	Synchronous and asynchronous replication	Mirrors data from a single primary site to two secondary sites simultaneously.	Disaster recovery scenarios depend on support from controlling software such as Geographically Dispersed Parallel Sysplex (GDPS) and Tivoli Storage Productivity Center for Replication
Metro/Global Mirror	Three-site, long distance disaster recovery replication	A backup site is maintained regardless of which one of the sites is lost.	Recovery point objective (RPO) might grow if bandwidth capability is exceeded.
Metro Mirror	Synchronous data copy at a distance	No data loss, rapid recovery time for distances up to 300 km.	Slight performance impact.
Global Copy	Continuous copy without data consistency	Nearly unlimited distance, suitable for data migration, only limited by network and channel extenders capabilities.	Copy is normally fuzzy but can be made consistent through synchronization.
Global Mirror	Asynchronous copy	Nearly unlimited distance, scalable, and low RPO. The RPO is the time needed to recover from a disaster; that is, the total system downtime.	RPO might grow when link bandwidth capability is exceeded.
z/OS Global Mirror	Asynchronous copy controlled by z Systems host software	Nearly unlimited distance, highly scalable, and very low RPO.	Additional host server hardware and software is required. The RPO might grow if bandwidth capability is exceeded or host performance might be impacted.

I/O Priority Manager

The performance group attribute associates the logical volume with a performance group object. Each performance group has an associated performance policy which determines how the I/O Priority Manager processes I/O operations for the logical volume.

Note: The default setting for this feature is "disabled" and must be enabled for use.

The I/O Priority Manager maintains statistics for the set of logical volumes in each performance group that can be queried. If management is performed for the performance policy, the I/O Priority Manager controls the I/O operations of all managed performance groups to achieve the goals of the associated performance policies. The performance group defaults to 0 if not specified. Table 19 lists performance groups that are predefined and have the associated performance policies:

Table 19. Performance groups and policies

Performance group ¹	Performance policy	Performance policy description
0	0	No management
1-5	1	Fixed block high priority
6-10	2	Fixed block medium priority
11-15	3	Fixed block low priority
16-18	0	No management
19	19	CKD high priority 1
20	20	CKD high priority 2
21	21	CKD high priority 3
22	22	CKD medium priority 1
23	23	CKD medium priority 2
24	24	CKD medium priority 3
25	25	CKD medium priority 4
26	26	CKD low priority 1
27	27	CKD low priority 2
28	28	CKD low priority 3
29	29	CKD low priority 4
30	30	CKD low priority 5
31	31	CKD low priority 6
Note: ¹ Performance group settings can be managed using DS CLI or the DS Storage Manager.		

Securing data

You can secure data with the encryption features that are supported by the DS8000 storage system. The DS8000 series supports data encryption by using IBM Full Disk Encryption (FDE) feature and key managers, such as IBM Security Key Lifecycle Manager.

Encryption technology has a number of considerations that are critical to understand to maintain the security and accessibility of encrypted data. For example, encryption must be enabled by feature code and configured to protect data in your environment. Encryption is not automatically activated because FDE drives are present.

It is important to understand how to manage IBM encrypted storage and comply with IBM encryption requirements. Failure to follow these requirements might cause a permanent encryption deadlock, which might result in the permanent loss of all key-server-managed encrypted data at all of your installations.

The DS8000 system tests access to the encryption key that is stored on each key server once every 8 hours. You can now initiate a request to test that a specific encryption group has access to an encryption key on a key server on demand. Status of successful and failed attempts of this access can also be monitored by using the DS CLI.

Chapter 4. Planning the physical configuration

Physical configuration planning is your responsibility. Your IBM representative can help you to plan for the physical configuration and to select features.

This section includes the following information:

- Explanations for available features that can be added to the physical configuration of your system model
- Feature codes to use when you order each feature
- Configuration rules and guidelines

Configuration controls

Indicator features control the physical configuration of the storage system.

These indicator features are for administrative use only. The indicator features ensure that each storage system (the base frame plus any expansion frames) has a valid configuration. There is no charge for these features.

Your storage system can include the following indicators:

Expansion-frame position indicators

Expansion-frame position indicators flag models that are attached to expansion frames. They also flag the position of each expansion frame within the storage system. For example, a position 1 indicator flags the expansion frame as the first expansion frame within the storage system.

Standby CoD indicators

Each model contains a Standby CoD indicator that indicates whether the storage system takes advantage of the Standby Capacity on Demand (Standby CoD) offering.

The warranty start date begins on the date that the CoD features are installed.

Administrative indicators

If applicable, models also include the following indicators:

- IBM / Openwave alliance
- IBM / EPIC attachment
- IBM systems, including System p, System x, IBM z Systems, and BlaceCenter
- IBM storage systems, including IBM System Storage ProtecTIER, IBM Storwize V7000, and IBM System Storage N series
- IBM SAN Volume Controller
- Linux
- VMware VAAI indicator
- Global Mirror

Determining physical configuration features

You must consider several guidelines for determining and then ordering the features that you require to customize your storage system. Determine the feature codes for the optional features you select and use those feature codes to complete your configuration.

Procedure

1. Calculate your overall storage needs. Consider any licensed functions, such as FlashCopy and Remote Mirror and Copy, that are used to ensure continuous data availability and to implement disaster-recovery requirements.

Note: If you are activating features for any of the licensed functions, such as Copy Services, all features must have the same capacity, including the operating environment license feature.

2. Determine the base and expansion models of which your storage system is to be comprised.
3. Determine the management console configuration that supports the storage system by using the following steps:
 - a. Order one internal management console for each storage system. The internal management console feature code must be ordered for the base model within the storage system.
 - b. Decide whether an external management console is to be installed for the storage system. Adding an external management console ensures that you maintain a highly available environment.
4. For each base and expansion model, determine the storage features that you need.
 - a. Select the drive set feature codes and determine the amount of each feature code that you must order for each model.
 - b. Select the storage enclosure feature codes and determine the amount that you must order to enclose the drive sets that you are ordering.
 - c. Select the disk cable feature codes and determine the amount that you need of each.
5. Determine the I/O adapter features that you need for your storage system.
 - a. Select the device, flash RAID, and host adapters feature codes to order, and choose a model to contain the adapters. All base models can contain adapters, but only the first attached expansion model can contain adapters.
 - b. For each model chosen to contain adapters, determine the number of each I/O enclosure feature codes that you must order.
 - c. Select the cables that you require to support the adapters.
6. Based on the disk storage and adapters that the base model and expansion models support, determine the appropriate processor memory feature code that is needed by each base model.
7. Decide which power features that you must order to support each model.
8. Review the other features and determine which feature codes to order.

Management console features

Management consoles are required features for your storage system configuration.

Customize your management consoles by specifying the following different features:

- An external management console and the required internal management console
- Management console external power cords

Internal and external management consoles

The management console is the focal point for configuration, Copy Services functions, remote support, and maintenance of your storage system.

The management console (also known as the Hardware Management Console or HMC) is a dedicated appliance that is physically located inside your storage system. It can proactively monitor the state of your storage system and notify you and IBM when service is required. It also can be connected to your network for centralized management of your storage system by using the IBM DS command-line interface (DS CLI) or storage management software through the IBM DS Open API. (The DS8000 Storage Management GUI cannot be started from the HMC.)

You can also use the DS CLI to control the remote access of your IBM service representative to the HMC.

An external management console is available as an optional feature. The external HMC is a redundant management console for environments with high-availability requirements. If you use Copy Services, a redundant management console configuration is especially important.

The internal management console is included with every base frame and is mounted in a pull-out tray for convenience and security. The external management console must be installed in an external 19-inch rack. This rack can be an IBM rack or a rack from another company. The rack must conform to the required specifications. When you order an external management console, the feature includes the hardware that is needed to install the management console in the rack.

Tip: To ensure that your IBM service representative can quickly and easily access an external management console, place the external management console rack within 15.2 m (50 ft) of the storage systems that are connected to it.

Notes:

1. To preserve console function, the external and the internal management consoles are not available as a general-purpose computing resource.
2. The external management console satisfies all applicable requirements of Section 508 of the Rehabilitation Act when assistive technology correctly interoperates with it.

Feature codes for management consoles

Use these feature codes to order up to two management consoles (MCs) for each storage system.

Table 20. Feature codes for the management console

Feature code	Description	Models
1120	Internal management console	A required feature that is installed in the 961 frame
1130	External management console	An optional feature that can be installed in an external IBM or a non-IBM rack

Management console external power cord

If you use an external management console, you must select an external power cord that is specific to your country, voltage, and amperage needs.

The power cord supplies external power to the external management console.

Feature codes for external management console power cords

Use these feature codes to order a power cord when you use an external management console.

Table 21. Feature codes for external management-console power cords

Feature code	Description	Country or region
1170	MC power cord standard rack	All
1171	MC power cord group 1	United States, Canada, Bahamas, Barbados, Bermuda, Bolivia, Brazil, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Guyana, Honduras, Jamaica, Japan, Japan (PDS), Mexico, Netherlands Antilles, Panama, Philippines, Saudi Arabia, Suriname, Taiwan, Trinidad, Venezuela
1172	MC power cord group 2 (250 V, 15 A)	Brazil

Configuration rules for management consoles

The management console is a dedicated appliance in your storage system that can proactively monitor the state of your storage system. You must order an internal management console each time that you order a base frame.

You can also order a second management console for your storage system. The second management console must be an external management console.

You must specify one keyboard feature with each management console that you order. Keyboard features specify the language and whether the keyboard is installed on an internal or external management console.

When you order an internal or external management console, the necessary Ethernet cables that attach it to the storage system are included.

Storage features

You must select the storage features that you want on your storage system.

The storage features are separated into the following categories:

- Drive-set features and storage-enclosure features
- Standby Capacity on Demand features
- Enclosure filler features
- Device adapter features

Storage enclosures and drives

DS8870 supports various storage enclosures and drive options.

Standard drive enclosures and drives

Standard drive enclosures and drives are required components of your storage system configuration.

Each standard drive enclosure feature contains two enclosures.

Each *drive set* feature contains 16 disk drives or flash drives (SSDs) and is installed with eight drives in each standard drive-enclosure pair.

Each *half-drive set* feature contains eight drives and is installed with four drives in each standard drive-enclosure pair.

The 3.5-inch storage enclosure slots are numbered left to right, and then top to bottom. The top row of drives is D01 - D04. The second row of drives is D05 - D08. The third row of drives is D09 - D12.

The 2.5-inch storage enclosure slots are numbered from left to right as slots D01 - D24. For full SFF (2.5-inch) drive sets, the first installation group populates D01 - D08 for both standard drive enclosures in the pair. The second installation group populates D09 - D16. The third installation group populates D17 - D24. For half-drive sets, the first installation group populates as D01 - D04 of both standard drive enclosures in the pair, and the second installation group populates as D05 - D08 in both standard drive enclosures, and so on.

Note: Storage enclosures are installed in the frame from the bottom up.

Table 22 and Table 23 provide information on the placement of full and half drive sets in the storage enclosure.

Table 22. Placement of full drive sets in the storage enclosure

Standard drive-enclosures type	Set 1	Set 2	Set 3
3.5 inch disk drives	D01 - D04	D05 - D08	D09 - D12
2.5 inch disk and flash drives	D01 - D08	D09 - D16	D17 - D24

Table 23. Placement of half drive sets in the storage enclosure

Standard drive-enclosures type	Set 1	Set 2	Set 3
3.5 inch disk drives	D01 - D04 (row 1)	D05 - D08 (row 2)	D09 - D12 (row 3)
2.5 inch disk and flash drives	D01 - D04	D05 - D08	D09 - D16

High-performance flash enclosures and flash cards

The DS8870 high performance flash enclosure provides significant performance improvements over prior generation SSDs and are the best choice for I/O intensive workloads.

The high-performance flash enclosure requires 1U of rack space and two PCIe slots to attach the storage. The flash enclosure is directly attached to the PCIe fabric, which increases bandwidth and transaction-processing capability.

Each flash enclosure feature is populated with 16 or 30 400-GB 1.8-inch flash-card set feature.

Each flash enclosure contains a pair of powerful redundant RAID controllers (known as *flash RAID adapters*) that are designed to unleash the performance capabilities of flash-based storage. Currently, the flash RAID adapters can be configured only for RAID 5 arrays.

The flash RAID adapters are each attached to a different I/O enclosure by using a x4 Gen2 PCIe cable. The cable is driven by a flash interface card that is plugged into an I/O enclosure slot.

DS8870 offers several types of configurations:

- Up to eight flash enclosures in an all-flash, single rack system with up to 96 TB of physical capacity. The DS8870 all-flash configuration provides twice as many I/O enclosures and up to twice the host adapters.
- Up to eight high-performance flash enclosures, four each in the base frame and first expansion frame, in addition to the existing standard configuration.

Standby CoD drive sets

You can use the IBM Standby Capacity on Demand (Standby CoD) feature by ordering Standby CoD disk-drive sets.

Capacity on demand (CoD) drive features are 2.5-inch or 3.5-inch disk-drive sets that are physically installed in standard drive enclosures but are not activated for use until more capacity is needed. Standby CoD disk-drive sets contain 8 or 16 disk drive of the same drive type, capacity, and speed (7,200, 10,000, or 15,000 RPM).

Note: The flash drives are not available as Standby CoD drives.

To activate Standby CoD drives (and exchange them for regular disk drives), you must order a feature exchange. Ordering the feature exchange results in the removal of the Standby CoD feature and the addition of the corresponding regular disk drive feature of the same capacity and speed. The transaction is invoiced at the differential price between the features that are removed and added.

When you initially order Standby CoD drive features, you must sign a Standby CoD agreement. Each subsequent order for Standby CoD features requires a supplement to the agreement.

Feature codes for drive sets

Use these feature codes to order sets of encryption disk drives, flash drives, flash cards, and encryption standby capacity on demand (CoD) disk drives for DS8870.

All drives that are installed in a standard drive enclosure pair or single flash enclosure must be of the same drive type, capacity, and speed.

The 3 TB and 4 TB disk drives are installed in standard drive-enclosure pairs for 3.5-inch disk drives (feature code 1244). The 1.8" 400 GB Flash cards are installed in the high-performance flash enclosure. All other disk drives are installed in standard drive-enclosure pairs for 2.5-inch disk drives (feature code 1241). Each standard drive-enclosure pair can support up to three disk-drive sets (up to 48 2.5-inch disk drives or 24 3.5-inch disk drives).

A maximum of six features is allowed for standby capacity on demand (CoD) disk-drive-sets features.

The flash drives (SSDs) can be installed only in standard drive enclosures for flash drives. Each flash-drive set includes 16 flash drives. The 400 GB flash drives are also available in a half set of eight flash drives. The minimum number of supported flash drives (RAID-5 only) is eight.

The flash cards can be installed only in high-performance flash enclosures. See Table 26 on page 94 for the feature codes. Each flash enclosure can contain 16 or 30 flash cards. Flash-card set A (16 flash cards) is a required feature for each flash enclosure. Flash-card set B (14 flash cards) is an optional feature and can be ordered to fill the flash enclosure with a maximum of 30 flash cards. All flash cards in a flash enclosure must be the same type and same capacity.

Use these feature codes to order operating-environment licenses for your storage system. An operating-environment license is required per TB unit and per value unit for every storage system (including the base frame and all physically attached expansion frames). The extent of IBM authorization that is acquired through the function-authorization feature codes must cover the physical capacity and the value units for each drive that is installed in the storage system, excluding Standby CoD capacity.

An activation license (feature code 1750) is required to activate encryption. A deactivation license (feature code 1754) can be used to deactivate encryption.

For Standby CoD features, the warranty start date begins when the CoD features are installed.

Table 24, Table 25 on page 94, Table 26 on page 94, and Table 27 on page 94 list the feature codes for encryption drive sets and standby CoD disk-drive sets based on drive size and speed.

Table 24. Feature codes for disk-drive sets

Feature code	Disk size	Drive type	Drives per set	Drive speed in RPM (K=1000)	Encryption drive	RAID support	Required value units
5108	146 GB	2.5-in. disk drives	16	15 K	Yes	5, 6, 10	4.8
5308	300 GB	2.5-in. disk drives	16	15 K	Yes	5, 6, 10	6.8
5708	600 GB	2.5-in. disk drives	16	10 K	Yes	5, 6, 10	11.5
5618	600 GB	2.5-in. disk drives	16	15 K	Yes	5, 6, 10	11.5
5758	900 GB	2.5-in. disk drives	16	10 K	Yes	5, 6, 10	16.0
5768	1.2 TB	2.5-in disk drives	16	10 K	Yes	5, 6, 10	20.0
5858	3 TB	3.5-in. NL disk drives	8	7.2 K	Yes	6, 10	13.5
5868 ¹	4 TB	3.5-in. NL disk drives	8	7.2 K	Yes	6, 10	16.2

Note:

1. Drives are full disk encryption (FDE) self-encrypting drive (SED) capable.

Table 25. Feature codes for flash-drive sets

Feature code	Disk size	Drive type	Drives per set	Drive speed in RPM (K=1000)	Encryption drive	RAID support	Required value units
6068	200 GB	2.5-in flash drives	16	N/A	Yes	5	21.6
6156	400 GB	2.5-in flash drives	8	N/A	Yes	5	18.2
6158	400 GB	2.5-in flash drives	16	N/A	Yes	5	36.4
6258	800 GB	2.5-in flash drives	16	N/A	Yes	5	64.0
6358	1.6 TB	2.5-in flash drives	16	N/A	Yes	5, 10 ¹	125

Note:

1. RAID-10 available by request price quotation.

Table 26. Feature codes for flash-card sets

Feature code	Disk size	Drive type	Drives per set	Drive speed in RPM (K=1000)	Encryption drive	RAID support	Required value units
1506 ^{1, 2}	400 GB	1.8-in. flash cards	16	N/A	Yes	5	36.4
1508 ^{2, 3}	400 GB	1.8-in. flash cards	14	N/A	Yes	5	31.8
1596 ⁴	400 GB	1.8-in. flash cards	16	N/A	Yes	5	36.4
1598 ^{4, 5}	400 GB	1.8-in. flash cards	14	N/A	Yes	5	31.8

Note:

1. Required for each high-performance flash enclosure (feature code 1500).
2. Licensed machine code (LMC) V7.3 or later is required.
3. Optional with feature code 1506. If feature code 1508 is not ordered, a storage filler set (feature code 1599) is required.
4. Licensed machine code (LMC) V7.4.1 or later is required.
5. Optional with feature code 1596. If feature code 1598 is not ordered, a storage filler set (feature code 1599) is required.

Table 27. Feature codes for standby CoD disk-drive sets

Feature code	Disk size	Drive type	Number of drives per set	Drive speed	Encryption drive	Required value units
5209	146 GB	2.5-in. disk drives	16	15 K	Yes	N/A
5309	300 GB	2.5-in. disk drives	16	15 K	Yes	N/A
5619	600 GB	2.5-in. disk drives	16	15 K	Yes	N/A
5709	600 GB	2.5-in. disk drives	16	10 K	Yes	N/A

Table 27. Feature codes for standby CoD disk-drive sets (continued)

Feature code	Disk size	Drive type	Number of drives per set	Drive speed	Encryption drive	Required value units
5759	900 GB	2.5-in. disk drives	16	10 K	Yes	N/A
5769	1.2 TB	2.5-in. disk drives	16	10 K	Yes	N/A
5859	3 TB	3.5-in. disk drives	8	7.2 K	Yes	N/A
5869	4 TB	3.5-in. disk drives	8	7.2 K	Yes	N/A

Feature codes for storage enclosures

Use these feature codes to order standard drive enclosures and high-performance flash enclosures for your storage system.

Table 28. Feature codes for storage enclosures

Feature code	Description	Models
1241	Standard drive-enclosure pair Note: This feature contains two filler sets in each enclosure. The enclosure pair supports the installation of one to three disk-drive set features.	961, 96E
1242	Standard drive-enclosure for 2.5-inch disk drives	961, 96E
1244	Standard drive-enclosure pair for 3.5-inch disk drives	961, 96E
1245	Standard drive-enclosure for 400 GB flash drives	961, 96E
1255	Standard drive-enclosure for 200 GB flash drives	961, 96E
1256	Standard drive-enclosure for 800 GB flash drives	961, 96E
1257	1.6 TB SSD enclosure indicator	961, 96E
1500	High-performance flash enclosure for flash cards	961

Storage-enclosure fillers

Storage-enclosure fillers fill empty drive slots in the storage enclosures. The fillers ensure sufficient airflow across populated storage.

For standard drive enclosures, one filler feature provides a set of 8 or 16 fillers. Two filler features are required if only one drive set feature is in the standard drive-enclosure pair. One filler feature is required if two drive-set features are in the standard drive-enclosure pair.

For high-performance flash enclosures, one filler feature provides a set of 14 fillers.

Feature codes for storage enclosure fillers

Use these feature codes to order filler sets for standard drive enclosures and flash enclosures.

Table 29. Feature codes for storage enclosures

Feature code	Description
2997	Filler set for 3.5-in. standard disk-drive enclosures; includes eight fillers
2998	Filler set for 2.5-in. standard disk-drive enclosures; includes eight fillers
2999	Filler set for 2.5-in. standard disk-drive enclosures; includes 16 fillers
1599	Filler set of flash enclosures; includes 14 fillers

Device adapters flash RAID adapters, and flash interface cards

Device adapters and flash interface cards provide the connection between storage devices and the internal processors and memory. Device adapters and flash RAID adapters perform the RAID management and control functions of the drives that are attached to them.

Each pair of device adapters or flash interface cards supports two independent paths to all of the drives that are served by the pair. Two paths connect to two different network fabrics to provide fault tolerance and to ensure availability. By using physical links, two read operations and two write operations can be performed simultaneously around the fabric.

Device adapters are ordered in pairs. For storage systems that use standard drive enclosures, the device adapters are installed in the I/O enclosure pairs, with one device adapter in each I/O enclosure of the pair. The device adapter pair connects to the standard drive enclosures by using 8 Gbps FC-AL. An I/O enclosure pair can support up to two device adapter pairs.

Flash RAID adapters are integrated as a pair in the high-performance flash enclosure feature. For storage systems that use high-performance flash enclosures, a pair of flash interface cards is installed in a pair of I/O enclosures to attach the high-performance flash enclosure. Each I/O enclosure can support up to two flash RAID adapter pairs.

Feature codes for device adapters and flash interface cards

Use these feature codes to order device adapters or flash interface cards for your storage system. Each feature includes two adapters.

Table 30. Feature codes for device adapters and flash interface cards

Feature code	Description	Models
3053	4-port, 8 Gbps device adapter pair	961, 96E
3054 ¹	Flash interface card pair	961, 96E
Note:		
1. Included with each high-performance flash enclosure (feature code 1500).		

Drive cables

You must order at least one drive cable set to connect the disk drives to the device adapters.

The disk drive cable feature provides you with a complete set of Fibre Channel cables to connect all the disk drives that are supported by the model to their appropriate device adapters.

Disk drive cable groups have the following configuration guidelines:

- The minimum number of disk-drive cable group features for each model is one.
- The disk-drive cable groups must be ordered as follows:
 - If the disk drives connect to device adapters within the same base frame, order disk drive cable group 1.
 - If the disk drives connect to device adapters within the same expansion frame, order disk drive cable group 2.
 - If the disk drives are in a second expansion frame (position 2 expansion frame), order disk drive cable group 4.
 - If the disk drives are in a third expansion frame, order disk drive cable group 5.

Feature codes for drive cables

Use these feature codes to order the cable groups for your storage system.

Table 31. Feature codes for drive cable

Feature code	Description	Connection Type
1246	Drive cable group 1	Connects the drives to the device adapters within the same base model 961.
1250	Drive cable group 1	(DS8870 business-class only) Connects the drives from a third expansion model 96E to the second expansion model 96E.
1247	Drive cable group 2	(Enterprise-class) Connects the drives to the device adapters in the first expansion model 96E (Business-class) Connects the drives from the first expansion model 96E to the base model 961.
1248	Drive cable group 4	Connects the drives from a second expansion model 96E to the base model 961 and first expansion model 96E.
1249	Drive cable group 5	(Enterprise-class) Connects the drives from a third expansion model 96E to a second expansion model 96E. (Business-class) Not applicable

Configuration rules for storage features

Use the following general configuration rules and ordering information to help you order storage features.

High-performance flash enclosures

Follow these configuration rules when you order storage features for storage systems with high-performance flash enclosures.

Flash enclosures

For systems with an All-Flash configuration, up to eight Flash enclosures are supported on a single frame. Four high-performance flash enclosures in

the base frame and four high-performance flash enclosures in the first expansion frame. For systems with an Enterprise Class or Business Class configuration, up to four flash enclosures are supported in the base frame.

For configurations with 16 GB system memory (feature 4311), flash enclosures are not supported. For configurations with 32 GB system memory (feature 4312), a maximum of two flash enclosures are supported.

Flash-card sets

Each high-performance flash enclosure requires a minimum of one 16 flash-card set.

Storage enclosure fillers

For high-performance flash enclosures, one filler feature provides a set of 14 fillers. One filler feature is required if only the optional 14 flash-card set is not ordered for the flash enclosure.

Flash RAID adapters

The flash RAID adapters are included in the high-performance flash enclosure feature and are not ordered separately.

Flash interface cards

Flash-interface cards come in pairs. One flash-interface-card pair is included with each high-performance flash enclosure feature and is not ordered separately.

Drive cables

One drive cable set is required to connect the high-performance flash enclosure to the flash-interface cards in the I/O enclosure.

Standard drive enclosures

Follow these configuration rules when you order storage features for storage systems with standard drive enclosures.

Standard drive enclosures

Storage enclosures are installed from the bottom to the top of each base or expansion frame. Depending on the number of drive sets that you order for the expansion frame, you might be required to fully populated the standard drive enclosure before you can order the next required drive sets.

Drive sets

Each standard high-density drive enclosure requires a minimum of eight flash drives, disk drives, or Standby CoD disk drives. The drive features that you order for the standard drive enclosure must be of the same type, capacity, and speed.

Each base frame requires a minimum of one drive set.

When you initially order Standby CoD drive features, you must sign a Standby CoD agreement. Each subsequent order for Standby CoD features requires a supplement to the agreement.

To activate Standby CoD drives (and exchange them for regular drives), you must order a feature exchange. Ordering the feature exchange results in the removal of the Standby CoD feature and the addition of the corresponding regular drive feature of the same type, capacity, and speed. The transaction is invoiced at the differential price between the features that are removed and added.

Storage enclosure fillers

One filler feature provides a set of 8 or 16 fillers. Two filler features are

required if only one drive-set feature is ordered for a standard drive-enclosure pair. One filler feature is required if two drive-set features are ordered for the standard drive-enclosure pair.

Device adapters

Device adapters are ordered in pairs. A minimum of one pair is required for each base frame.

Enterprise configuration device adapter pairs must be ordered as follows:

- For configurations with 16 GB or 32 GB system memory (feature 4311 or 4312), up to two device-adapter pairs are supported.
- For configurations with 64 GB system memory, up to four device-adapter pairs are supported.
- For configurations with 128 GB system memory, up to eight device-adapter pairs are supported.

Drive cables

At least one drive cable set is required to connect the disk drives to the device adapters.

The disk-drive cable groups must be ordered as follows:

- If the drives connect to device adapters within the same base frame, order drive-cable group 1.
- If the drives connect to device adapters within the same expansion frame, order drive-cable group 2.
- If the drives are in a second expansion frame (position 2 expansion frame), order drive-cable group 4.
- If the drives are in a third expansion frame, order drive-cable group 5.

Physical and effective capacity

Use the following information to calculate the physical and effective capacity of a storage system.

To calculate the total physical capacity of a storage system, multiply each drive-set feature by its total physical capacity and sum the values. For the standard drive enclosures, a full drive-set feature consists of 16 identical disk drives with the same drive type, capacity, and speed. For high-performance flash enclosures, there are 2 drive sets, one with 16 identical flash cards, and the other with 14 identical flash cards.

The logical configuration of your storage affects the effective capacity of the drive set.

Specifically, effective capacities vary depending on the following configurations:

- Data format

Physical capacity can be logically configured as fixed block (FB) or count key data (CKD). Data that is accessed by open systems hosts or Linux on IBM z Systems that support Fibre Channel protocol must be logically configured as FB. Data that is accessed by IBM z Systems hosts with z/OS or z/VM[®] must be configured as CKD.

- RAID ranks and configurations

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

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

Each RAID rank is divided into equal-sized segments that are known as *extents*. All extents are approximately 1 GB. However, CKD extents are slightly smaller than FB extents.

RAID capacities for DS8870

Use the following information to calculate the physical and effective capacity for DS8870.

RAID 5 array capacities

The following table lists the RAID 5 array capacities for fully populated storage enclosures.

Drive size and type	Total physical capacity (GB) per drive set ^{1, 2}	Fixed block (FB) or count key data (CKD)	Effective capacity in GB (number of extents) ^{3, 4}	
			Rank with RAID 5 array	
			6 + P	7 + P
146 GB disk drives	2,336	FB	794.569 (740)	933.082 (869)
		CKD	793.650 (829)	931.509 (973)
200 GB 2.5-in. flash drives (SSD)	3,200	FB	1023.276 (953)	1198.296 (1116)
		CKD	1021.501 (1067)	1197.655 (1251)
300 GB disk drives	4,800	FB	1655.710 (1542)	1935.957 (1803)
		CKD	1653.357 (1727)	1933.863 (2020)
400 2.5-in. GB flash drives (SSD)	6,400	FB	2277.406 (2121)	2660.732 (2478)
		CKD	2274.683 (2376)	2658.583 (2777)
400 GB 1.8-in. flash cards ²	6,400	FB	2279.55 (2123)	n/a
		CKD	2249.60 (2378)	n/a
600 GB disk drives	9,600	FB	3372.623 (3141)	3936.338 (3666)
		CKD	3367.986 (3518)	3931.870 (4107)
800 GB 2.5-in. flash drive (SSD)	12,800	FB	4,594.541 (4279)	5,361.193 (4993)
		CKD	4,587.660 (4792)	5,354.504 (5593)

Drive size and type	Total physical capacity (GB) per drive set ^{1, 2}	Fixed block (FB) or count key data (CKD)	Effective capacity in GB (number of extents) ^{3, 4}	
			Rank with RAID 5 array	
			6 + P	7 + P
900 GB disk drives	14,400	FB	5077.725 (4729)	5924.907 (5518)
		CKD	5071.126 (5297)	5917.430 (6181)
1.2 TB disk drives	19,200	FB	6782.827 (6317)	7912.404 (7369)
		CKD	6774.266 (7076)	7902.991 (8255)
1.6 TB flash drive (SSD)	25,600	FB	9226.663 (8593)	10761.04 (10,022)
		CKD	9105.303 (9625)	10620.80 (11,227)
3 TB disk drives (Nearline)	48,000	FB	17,015.587 (15,847)	19,838.454 (18,476)
		CKD	16,992.149 (17,749)	19,816.355 (20,699)
4 TB disk drives (Nearline)	64,000	FB	22,701.050 (21,142)	26,464.515 (24,647)
		CKD	22,669.282 (23,679)	26,434.571 (27,612)

Notes:

- Disk-drive and flash-drive sets contain 16 disk drives. Half-drive sets contain 8 disk drives.
- The high-performance flash enclosure is populated with 30 flash cards (feature code 1506, 16 flash cards and feature code 1508, 14 flash cards) or with 16 flash cards (feature code 1506) and 14 filler sets (feature code 1599).
- Physical capacities are in decimal gigabytes (GB) and terabytes (TB). One decimal GB is 1,000,000,000 bytes. One decimal TB is 1,000,000,000,000 bytes.
- Rank capacities are slightly smaller on DS8870 as compared to the similar DS8800. This reduction of extents must be planned for when you move or migrate data to the DS8870.

RAID 6 array capacities

The following table lists the RAID 6 array capacities for fully populated storage enclosures.

Drive size and type	Total physical capacity (GB) per drive set ^{1, 2}	Fixed block (FB) or count key data (CKD)	Effective capacity in GB (number of extents) ^{2, 3}	
			Rank with RAID 6 array	
			5 + P + Q	6 + P + Q
146 GB disk drives	2,336	FB	639.950 (596)	777.389 (724)
		CKD	639.515 (668)	777.375 (812)

Drive size and type	Total physical capacity (GB) per drive set ^{1, 2}	Fixed block (FB) or count key data (CKD)	Effective capacity in GB (number of extents) ^{2, 3}	
			Rank with RAID 6 array	
			5 + P + Q	6 + P + Q
300 GB disk drives	4,800	FB	1341.104 (1249)	1621.350 (1510)
		CKD	1340.301 (1400)	1618.893 (1691)
600 GB disk drives	9,600	FB	2738.042 (2550)	3301.756 (3075)
		CKD	2736.129 (2858)	3298.099 (3445)
900 GB disk drives	14,400	FB	4125.316 (3842)	4971.425 (4630)
		CKD	4122.384 (4306)	4966.774 (5188)
1.2 TB disk drives	19,200	FB	5513.664 (5135)	6642.167 (6186)
		CKD	5509.596 (5755)	6634.491 (6930)
3 TB disk drives (Nearline)	48,000	FB	13,840.532 (12,890)	16,662.326 (15,518)
		CKD	13,831.910 (14,448)	16,644.628 (17,386)
4 TB disk drives (Nearline)	64,000	FB	18,467.286 (17,199)	22,228.603 (20,702)
		CKD	18,454.992 (19,277)	22,205.921 (23,195)
Notes: 1. Disk-drive and flash-drive sets contain 16 disk drives. Half-drive sets contain 8 disk drives. 2. Physical capacities are in decimal gigabytes (GB) and terabytes (TB). One decimal GB is 1,000,000,000 bytes. One decimal TB is 1,000,000,000,000 bytes. 3. Rank capacities are slightly smaller on DS8870 as compared to the similar DS8800. This reduction of extents must be planned for when you move or migrate data to the DS8870.				

RAID 10 array capacities

The following table lists the RAID 10 array capacities for fully populated storage enclosures.

Drive size and type	Total physical capacity (GB) per drive set ¹	Fixed block (FB) or count key data (CKD)	Effective capacity in GB (number of extents) ^{2, 3}	
			Rank with RAID 10 array	
			3 + 3	4 + 4
146 GB disk drives	2,336	FB	379.031 (353)	517.544 (482)
		CKD	378.156 (395)	516.973 (540)
200 GB 2.5-in. flash drive (SSD)	3,200	FB	492.847 (459)	670.015 (624)
		CKD	492.082 (514)	669.193 (699)
300 GB disk drives	4,800	FB	809.601 (754)	1091.995 (1017)
		CKD	808.010 (844)	1090.431 (1139)
400 2.5-in. GB flash drives (SSD)	6,400	FB	1120.986 (1044)	1507.534 (1404)
		CKD	1119.152 (1169)	1504.010 (1571)
600 GB disk drives	9,600	FB	1667.521 (1553)	2236.604 (2083)
		CKD	1665.803 (1740)	2232.559 (2332)
800 GB 2.5-in. flash drive (SSD)	12,800	FB	2278.480 (2122)	3051.574 (2842)
		CKD	2275.640 (2377)	3046.313 (3182)
900 GB disk drives	14,400	FB	2520.072 (2347)	3374.771 (3143)
		CKD	2516.894 (2629)	3367.986 (3518)
1.2 TB disk drives	19,200	FB	3373.697 (3142)	4511.863 (4202)
		CKD	3368.943 (3519)	4503.412 (4704)
1.6 TB flash drive (SSD)	25,600	FB	4595.615 (4280)	6141.803 (5720)
		CKD	4534.204 (4793)	6058.219 (6404)
3 TB disk drives (Nearline)	48,000	FB	8,490.077 (7907)	11,337.640 (10,559)
		CKD	8477.406 (8855)	11,315.973 (11,820)
4 TB disk drives (Nearline)	64,000	FB	11,332.271 (10,554)	15,129.022 (14,090)
		CKD	11,315.973 (11,820)	15,100.409 (15,773)

Drive size and type	Total physical capacity (GB) per drive set ¹	Fixed block (FB) or count key data (CKD)	Effective capacity in GB (number of extents) ^{2, 3}	
			Rank with RAID 10 array	
			3 + 3	4 + 4
Notes: 1. Disk-drive and flash-drive sets contain 16 disk drives. Half-drive sets contain 8 disk drives. 2. Physical capacities are in decimal gigabytes (GB) and terabytes (TB). One decimal GB is 1,000,000,000 bytes. One decimal TB is 1,000,000,000,000 bytes. 3. Rank capacities are slightly smaller on DS8870 as compared to the similar DS8800. This reduction of extents must be planned for when you move or migrate data to the DS8870.				

I/O adapter features

You must select the I/O adapter features that you want for your storage system.

The I/O adapter features are separated into the following categories:

- I/O enclosures
- Device adapters
- Flash interface cards
- Host adapters
- Host adapters Fibre Channel cables

I/O enclosures

I/O enclosures are required for your storage system configuration.

The I/O enclosures hold the I/O adapters and provide connectivity between the I/O adapters and the storage processors. I/O enclosures are ordered and installed in pairs.

The I/O adapters in the I/O enclosures can be either device, flash RAID, or host adapters. Each I/O enclosure pair can support up to four device adapters (two pairs), four flash RAID adapters, and four host adapters.

Feature codes for I/O enclosures

Use this feature code to order I/O enclosures for your storage system.

The I/O enclosure feature includes two I/O enclosures. This feature supports up to two device adapter pairs, two flash interface card pairs, and up to four host adapters.

Table 32. Feature codes for I/O enclosures

Feature code	Description
1301	I/O enclosure pair

Feature codes for I/O cables

Use these feature codes to order the I/O cables for your storage system.

Table 33. Feature codes for PCIe cables

Feature Code	Cable Group	Description	Models
1320	PCIe cable group 1	Connects device and host adapters in an I/O enclosure pair to the processor.	961
1321	PCIe cable group 2	Connects device and host adapters in I/O enclosure pairs to the processor.	961
1322	PCIe cable group 3	Connects device and host adapters in an I/O enclosure pair to the processor.	96E

Fibre Channel (SCSI-FCP and FICON) host adapters and cables

You can order Fibre Channel host adapters for your storage-system configuration.

The Fibre Channel host adapters enable the storage system to attach to Fibre Channel (SCSI-FCP) and FICON servers, and SAN fabric components. They are also used for remote mirror and copy control paths between DS8000 series storage systems, or between a DS8000 series storage system and a DS6000 series or a 2105 (model 800 or 750) storage system. Fibre Channel host adapters are installed in an I/O enclosure.

Adapters are either 4-port 4 Gbps, 4-port or 8-port 8 Gbps, or 4-port 16 Gbps. Each adapter type supports 4, 8, and 16 Gbps full-duplex Fibre Channel data transfer speeds.

Supported protocols include the following types:

- SCSI-FCP upper layer protocol (ULP) on point-to-point, fabric, and arbitrated loop (private loop) topologies.

Note: The 16Gbps adapter does not support arbitrated loop topology at any speed.

- FICON ULP on point-to-point and fabric topologies.

Notes:

1. SCSI-FCP and FICON are supported simultaneously on the same adapter, but not on the same port.
2. For highest availability, ensure that you add adapters in pairs.

A Fibre Channel cable is required to attach each Fibre Channel adapter port to a server or fabric component port. The Fibre Channel cables can be 50 or 9 micron, OM3 or higher fiber graded, single or multimode cables.

Feature codes for Fibre Channel host adapters

Use these feature codes to order Fibre Channel host adapters for your storage system.

A maximum of four Fibre Channel host adapters can be ordered with 2-core processor license (feature code 4411). A maximum of 16 Fibre Channel host

adapters can be ordered for DS8870 All Flash systems with 256 GB or higher processor memory (feature codes 4315 - 4317). All other configuration support a maximum of 16 Fibre Channel host adapters.

Table 34. Feature codes for Fibre Channel host adapters

Feature code	Description	Receptacle type
3153	4-port, 8 Gbps shortwave FCP and FICON host adapter, PCIe	LC
3157	8-port, 8 Gbps shortwave FCP and FICON host adapter, PCIe	LC
3253	4-port, 8 Gbps longwave FCP and FICON host adapter, PCIe	LC
3257	8-port, 8 Gbps longwave FCP and FICON host adapter, PCIe	LC
3353 ₁	4-port, 16 Gbps shortwave FCP and FICON host adapter, PCIe	LC
3453 ₁	4-port, 16 Gbps longwave FCP and FICON host adapter, PCIe	LC
<p>1. If you are replacing an existing host adapter with a higher speed host adapter, once the exchange process has started, do not make any changes to the host configuration until the replacement is complete. Port topology is restored during the exchange process, and the host configuration appears when the host adapter is installed successfully.</p>		

Feature codes for Fibre Channel cables

Use these feature codes to order Fibre Channel cables to connect Fibre Channel host adapters to your storage system. Take note of the distance capabilities for cable types.

Table 35. Feature codes for Fibre Channel cables

Feature code	Cable type	Cable length	Compatible Fibre Channel host adapter features
1410	50 micron OM3 or higher Fibre Channel cable, multimode	40 m (131 ft)	<ul style="list-style-type: none"> Shortwave Fibre Channel or FICON host adapters (feature codes 3153, 3157, 3453)
1411	50 micron OM3 or higher Fibre Channel cable, multimode	31 m (102 ft)	
1412	50 micron OM3 or higher Fibre Channel conversion cable, multimode	2 m (6.5 ft)	
1420	9 micron OS1 or higher Fibre Channel cable, single mode	31 m (102 ft)	<ul style="list-style-type: none"> Longwave Fibre Channel or FICON adapters (feature codes 3253, 3257, 3353).
1421	9 micron OS1 or higher Fibre Channel cable, single mode	31 m (102 ft)	
1422	9 micron OS1 or higher Fibre Channel conversion cable, single mode	2 m (6.5 ft)	

Table 36. Multimode cabling limits.

Fibre cable type	Distance limits relative to Gbps			
	2 Gbps	4 Gbps	8 Gbps	16 Gbps
OM1 (62.5 micron)	150 m	70 m	Not recommended	Not recommended
OM2 (50 micron)	300 m	150 m	Not recommended	35 m
OM3 (50 micron)	500 m	380 m	150 m	100 m
OM4 (50 micron)	500 m	400 m	190 m	125 m

Configuration rules for I/O adapter features

To order I/O adapter features, you must follow specific configuration rules.

The following configuration rules affect I/O adapter features:

- Configuration rules for I/O enclosures and adapters
- Configuration rules for host adapters and host adapter cables

Configuration rules for I/O enclosures, adapters, and cards

Use these configuration rules and ordering information to help you order I/O enclosures, adapters, and cards.

Use the following tables to determine the number of I/O enclosures, device adapters, and flash interface features that you need in various storage configurations. To use the table, find the rows that contain the type of storage system that you are configuring. Then, find the row that represents the number of storage enclosures that are installed in that storage system. Use the remaining columns to find the number of I/O enclosure, device adapters, and flash interface cards that you need in the storage system.

Table 37. Required I/O enclosures and flash interface cards for All Flash configurations

Processor type	Storage frame	Flash enclosure features	Required flash interface card features (3054) ¹	I/O enclosure pair features (1301)
8-core or 16-core	Base frame	1	1	4
		2	2	4
		3	3	4
		4	4	4
		5	5	4
		6	6	4
		7	7	4
		8	8	4

Notes:

1. Each flash interface card feature code represents one flash-interface-card pair.

Table 38. Required I/O enclosures and device adapters Enterprise Class configurations

Processor type	Storage frame	Standard drive enclosure features (1241) ¹	Required device adapter features (3053) ²	Flash enclosure features (1500) ³	Required flash interface card features (3054) ⁴	Required I/O enclosure features (1301) ⁵
2-core	Base frame	1	1	0 - 2	0 - 2	1
		2	2	0 - 2	0 - 2	1
		3	2	0 - 2	0 - 2	1
4-core, 8-core, or 16-core	Base frame	1	1	0 - 4	0 - 4	2
		2	2	0 - 4	0 - 4	2
		3	3	0 - 4	0 - 4	2
		4-5	4	0 - 4	0 - 4	2
8-core or 16-core	Expansion frame	1	1	0 - 4	0 - 4	2
		2	2	0 - 4	0 - 4	2
		3	3	0 - 4	0 - 4	2
		4 - 7	4	0 - 4	0 - 4	2
8-core or 16-core	Second or third expansion frame	1 - 10	n/a	n/a	n/a	n/a

Notes:

1. Each storage enclosure feature represents one storage enclosure pair.
2. Each device adapter feature code represents one device adapter pair. The maximum quantity is two device adapter features for each I/O enclosure feature in the storage system.
3. Each flash interface card feature code represents one flash-interface-card pair. The maximum quantity is two flash-interface-card features for each I/O enclosure feature in the storage system.
4. Each flash enclosure feature represents a single enclosure, and requires a minimum of 32 GB memory.
5. Each I/O enclosure feature represents one I/O enclosure pair. An I/O enclosure pair can support up to two device adapter pairs and two flash interface card pairs.

Table 39. Required I/O enclosures and device adapters for Business Class configurations

Processor type	Storage frame	Standard drive enclosure features (1241)	Required device adapter features (3053)	Flash enclosure features (1500) ⁴	Required flash interface card features (3054) ⁵	Required I/O enclosure features (1301) ^{6,7}
2-core	Base frame	1 - 3	1	0 - 2	0 - 2	1
4-core, 8-core, or 16-core	Base frame	1 - 5	2	0 - 4	0 - 4	2
		1 - 5	2	0 - 4	0 - 4	2
8-core or 16-core	First expansion frame	0 - 3	2 ¹	0 - 4	0 - 4	2 ²
		4 - 7	3 ¹	0 - 4	0 - 4	2 ³

Table 39. Required I/O enclosures and device adapters for Business Class configurations (continued)

Processor type	Storage frame	Standard drive enclosure features (1241)	Required device adapter features (3053)	Flash enclosure features (1500) ⁴	Required flash interface card features (3054) ⁵	Required I/O enclosure features (1301) ^{6, 7}
8-core or 16-core	Second expansion frame	1 - 4	4 ¹	n/a	n/a	2 ²
		5 - 8	5 ¹	n/a	n/a	2 ³
		9 - 10	6 ¹	n/a	n/a	2 ³
Notes: <ol style="list-style-type: none"> The required device adapter is based on accumulative standard drive-enclosure features. The required I/O enclosure is installed in the base frame. The required I/O enclosure is installed in the first expansion frame. Each high-performance flash interface card feature code represents one flash-interface-card pair. The maximum quantity is two flash-interface-card features for each I/O enclosure feature in the storage system. Each high-performance flash enclosure feature represents a single enclosure, and requires a minimum of 32 GB memory. If you order more than one high-performance flash enclosure, you might need more than 1 I/O enclosure pair depending on positions of the flash enclosure. If position 1 and 2 are used, two I/O enclosure pairs are required. If position 1 and 3 are used, then one I/O enclosure is required. Each I/O enclosure feature represents one I/O enclosure pair. An I/O enclosure pair can support up to two device adapter pairs and two flash interface card pairs. 						

Configuration rules for host adapters

Use the following configuration rules and ordering information to help you order host adapters.

When you configure your storage system, consider the following issues when you order the host adapters:

- What are the minimum and maximum numbers of host adapters that I can install?
- How can I balance the host adapters across the storage system to help ensure optimum performance?
- What host adapter configurations help ensure high availability of my data?
- How many and what type of cables do I need to order to support the host adapters?

In addition, consider the following host adapter guidelines.

- You can include a combination of Fibre Channel host adapters in one I/O enclosure.
- Feature conversions are available to exchange installed adapters when new adapters of a different type are purchased.

Maximum and minimum configurations

The following table lists the minimum and maximum host adapter features for the base frame (model 961).

Table 40. Minimum and maximum host adapter features for the base frame

Storage system type	Storage system configuration	Minimum number of host adapter features for the base frame	Maximum number of host adapter features for the storage system ¹
2-core	Base frame	2	4
4-core	Base frame	2	8
8-core or 16-core	Base frame + 1-3 expansion frame	2	16

Note: ¹ For Enterprise Class and Business Class configurations, the maximum number of host adapters for any one frame cannot exceed 8 with 8 Gbps host adapters. You can add host adapters only to the base frame (model 961) and the first expansion frame (model 96E).

For All Flash configurations, the maximum number of host adapters for any one frame cannot exceed 16 8-Gbps host adapters in the base frame.

Configuring for highest availability

After you meet the initial minimum order requirement, you can order one host adapter at a time. However, it is recommended that you add host adapters (of the same type) in pairs.

For optimum performance, it is important that you aggregate the bandwidth across all the hardware by installing host adapters evenly across all available IO enclosures.

Notes:

- Although one multiport adapter can provide redundant pathing, keep in mind that any host requires access to data through a minimum of two separate host adapters. To maintain access to data, in the event of a host adapter failure or an I/O enclosure failure, the host adapters must be in different I/O enclosures.
- If an IBM service representative moves existing host adapters from one slot to another, you can configure the host ports on your storage system by using the IBM DS Storage Manager or the DS CLI.

Ordering host adapter cables

For each host adapter, you must provide the appropriate fiber-optic cables. Typically, to connect Fibre Channel host adapters to a server or fabric port, provide the following cables:

- For shortwave Fibre Channel host adapters, provide a 50-micron multimode OM3 or higher fiber-optic cable that ends in an LC connector.
- For longwave Fibre Channel host adapters, provide a 9-micron single mode OS1 or higher fiber-optic cable that ends in an LC connector.

These fiber-optic cables are available for order from IBM.

IBM Global Services Networking Services can assist with any unique cabling and installation requirements.

Processor complex features

These features specify the number and type of core processors in the processor complex. All base frames (model 961) contain two processor enclosures (POWER7+ servers) that contain the processors and memory that drives all functions in the storage system.

Feature codes for processor licenses

Use these processor-license feature codes to plan for and order processor memory for your storage system. You can order only one processor license per system.

POWER7+ processor modules require licensed machine code (LMC) V7.2 or later.

Expansion racks (model 96E) require the model 961 with 8-core or 16-core processors features.

Table 41. Feature codes for processor licenses

Feature code	Description	Corequisite feature code for memory
4411	2-core POWER7+ processor feature	4311 or 4312
4412	4-core POWER7+ processor feature	4313
4413	8-core POWER7+ processor feature	4314 or 4315
4414	16-core POWER7+ processor feature	4316 or 4317

Processor memory features

These features specify the amount of memory that you need depending on the processors in the storage system.

Feature codes for system memory

Use these feature codes to order system memory for your storage system.

Note: Memory is not the same as cache. The amount of cache is less than the amount of available memory. See the DS8000 Storage Management GUI

Table 42. Feature codes for system memory

Feature code	Description	Model
4311 ¹	16 GB system memory	961 (2-core)
4312 ¹	32 GB system memory	961 (2-core)
4313 ²	64 GB system memory	961 (4-core)
4314 ³	128 GB system memory	961 (8-core)
4315 ³	256 GB system memory	961 (8-core)
4316 ⁴	512 GB system memory	961 (16-core)
4317 ⁴	1 TB system memory	961 (16-core)

Table 42. Feature codes for system memory (continued)

Feature code	Description	Model
Notes:		
1. Feature codes 4311 and 4312 require 2-core processor license feature code 4411.		
2. Feature code 4313 requires 4-core processor license feature code 4412.		
3. Feature codes 4314 and 4315 require 8-core processor license feature code 4413.		
4. Feature codes 4316 and 4317 require 16-core processor license feature code 4414.		

Configuration rules for system memory

Use the configuration rules and ordering information to help you select and order system memory for your storage system.

You must order one system memory feature for the configuration of each base frame.

Model 961 2-core configuration, feature code 4311 and 4312

You can select 16 - 32 GB of system memory.

Note: The 16 GB (feature code 4211) and 32 GB (feature code 4212) system memory options are available only for the 2-core configuration feature.

Model 961 4-core configuration

Offers 64 GB of system memory.

Model 961 8-core configuration

You can select 128 - 256 GB of system memory.

Model 961 16-core configuration

You can select from 512 GB to 1 TB of system memory.

Power features

You must specify the power features to include on your storage system.

The power features are separated into the following categories:

- Power cords
- Input voltage
- DC-UPS (direct current uninterruptible power supply)

For the DS8870, base frame (models 961) and expansion frame (model 96E), the DC-UPS is included in your order.

Power cords

A pair of power cords (also known as *power cables*) is required for each base or expansion frame.

The DS8000 series has redundant primary power supplies. For redundancy, ensure that each power cord to the frame is supplied from an independent power source.

Feature codes for power cords

Use these feature codes to order power cords for DS8870 base or expansion racks. Each feature code includes two power cords. Ensure that you meet the requirements for each power cord and connector type that you order.

Important: IBM Safety requires a minimum of one IBM safety-approved ladder (feature code 1101) to be available at each installation site when overhead cabling (feature codes 1072, 1073, 1083, 1084) is used and when the maximum height of the overhead power source is 10 ft from the ground level. This ladder is a requirement for storage-system installation and service.

Table 43. Feature codes for power cords

Feature code	Power cord type	Wire gauge
1061	Single-phase power cord, 200-240V, 60A, 3-pin connector HBL360C6W, Pin and Sleeve Connector, IEC 309, 2P3W HBL360R6W, AC Receptacle, IEC 60309, 2P3W	10 mm ² (6awg)
1068	Single-phase power cord, 200-240V, 63A, no connector	10 mm ² (6awg)
1072	Top exit single-phase power cord, 200-240V, 60A, 3-pin connector	10 mm ² (6awg)
1073	Top exit single-phase power cord, 200-240V, 63A, no connector	10 mm ² (6awg)
1080	Three-phase power cord, high voltage (five-wire 3P+N+G) 380-415V (nominal), 30A, IEC 60309, 5-pin customer connector (available in North America) HBL530C6V02, Pin and Sleeve Connector, IEC 60309, 4P5W HBL530R6V02, AC Receptacle, IEC 60309, 4P5W	6 mm ² (10awg)
1081	Three-phase high voltage (five-wire 3P+N+G), 380-415V, 32A, no customer connector provided (available in Europe, Middle East, and Asia/Pacific)	6 mm ² (10awg)
1082	Three-phase power cord, low voltage, 200V-240V, 60A, 4-pin connector (available in United States, Canada, Latin America, Japan, and Asia/Pacific) HBL460C9W, Pin and Sleeve Connector, IEC 309, 3P4W HBL460R9W, AC Receptacle, IEC 60309, 3P4W	10 mm ² (6awg)
1083	Top exit power cord; Three-phase high voltage (five-wire 3P+N+G), 380-415V, 32A, no customer connector provided (available in Europe, Middle East, and Asia/Pacific)	6 mm ² (10awg)

Table 43. Feature codes for power cords (continued)

Feature code	Power cord type	Wire gauge
1084	<p>Top exit power cord; Three-phase low voltage (four-wire 3P+G), 200-240V, 60A, IEC 60309 4-pin customer connector (available in United States, Canada, Latin America, Japan, and Asia/Pacific)</p> <p>HBL460C9W, Pin and Sleeve Connector, IEC 309, 3P4W</p> <p>HBL460R9W, AC Receptacle, IEC 60309, 3P4W</p>	10 mm ² (6awg)
1085	<p>Top exit power cord; Three-phase high voltage (five-wire 3P+N+G), 380-415V, 30A (nominal), IEC 60309 5-pin customer connector (available in North America)</p> <p>HBL530C6V02, Pin and Sleeve Connector, IEC 60309, 4P5W</p> <p>HBL530R6V02, AC Receptacle, IEC 60309, 4P5W</p>	6 mm ² (10awg)

Input voltage

The DC-UPS distributes full wave, rectified power that ranges from 200 V ac to 240 V ac.

Direct-current uninterruptible-power supply

Each frame includes two direct-current uninterruptible-power supplies (DC-UPSs). Each DC-UPS can include one or two battery-service-module sets depending on the configuration ordered.

The DC-UPSs with integrated battery-service-module sets provides the ability to tolerate a power line disturbance. Loss of power to the frame up to 4 seconds is tolerated without a separate feature code. With the extended power line disturbance (ePLD) feature, loss of power for 50 seconds is tolerated without interruption of service.

The DC-UPS monitors its own alternating current (ac) input. Each DC-UPS rectifies and distributes the input ac. If a single DC-UPS in a frame loses ac input power, that DC-UPS receives and distributes rectified ac from the partner DC-UPS in that frame. If both DC-UPSs in that frame lose ac-input power, the DC-UPSs go “on battery.” If ac input is not restored within four seconds minimum (or 50 seconds with ePLD), the storage system commences shutdown.

Activation and Recovery for system failure: If both power cords lose ac input, the DC-UPS senses that both partner power and local power is running on batteries. Both stay on battery power and provide status to the rack power control (RPC), which initiates a recovery process.

Feature codes for battery service modules

Use this feature code to order battery service modules for your base and expansion racks.

Table 44. Feature code for the battery service module

Feature code	Description	Models
1051	Battery service module	All

Configuration rules for power features

Ensure that you are familiar with the configuration rules and feature codes before you order power features.

When you order power cord features, the following rules apply:

- You must order a minimum of one power cord feature for each frame. Each feature code represents a pair of power cords (two cords).
- You must select the power cord that is appropriate to the input voltage and geographic region where the storage system is located.

If the optional extended power line disturbance (ePLD) option is needed, you must order feature code 1055 for each base frame (model 961) and expansion frame (model 96E). The ePLD option protects the storage system from a power-line disturbance for 50 seconds. The following table lists the quantity of battery assembly features or extended PLD features (1055) you must order.

Table 45. Required quantity of battery assemblies

Model 961 and 96E with feature code 1055 or 1301	Model 961 and 96E without feature code 1055	Model 961 and 96E with feature code 1055
2 each	1 each	2 each

Other configuration features

Features are available for shipping and setting up the storage system.

You can select shipping and setup options for the storage system. The following list identifies optional feature codes that you can specify to customize or to receive your storage system.

- Extended power line disturbance (ePLD) option
- Remote IBM z Systems power control option
- Earthquake Resistance Kit option
- BSMI certificate (Taiwan)
- Shipping weight reduction option

Extended power line disturbance

The extended power line disturbance (ePLD) option (feature code 1055) gives the storage system the ability to tolerate power line disturbance for 50 seconds, rather than 4 seconds without the ePLD feature. This feature is optional for your storage system configuration.

If the ePLD option is not ordered, one battery service module (feature code 1051) per DC-UPS are required for each base and expansion frame. If the ePLD option is ordered, two battery service modules per DC-UPS are required for each base and expansion frame.

Feature code for extended power-line disturbance

Use this feature code to order the extended power-line disturbance (ePLD) feature for your storage system.

Table 46. Feature code for the ePLD

Feature code	Description	Models
1055	Extended power line disturbance	All

Remote zSeries power control feature

The optional remote zSeries power control feature adds a logic card that is used by one or more attached z Systems hosts to control the power-on and power-off sequences for your storage system.

When you use this feature, you must specify the **zSeries power control** setting in the DS8000 Storage Management GUI.

This feature includes the cables necessary to connect the logic card.

Feature code for remote zSeries power control

Use this feature code to order the remote zSeries power control feature for your storage system.

Table 47. Feature code for remote zSeries power control

Feature code	Description
1000	Remote zSeries power control

BSMI certificate (Taiwan)

The BSMI certificate for Taiwan option provides the required Bureau of Standards, Metrology, and Inspection (BSMI) ISO 9001 certification documents for storage system shipments to Taiwan.

If the storage system that you order is shipped to Taiwan, you must order this option for each frame that is shipped.

Feature code for BSMI certification documents (Taiwan)

Use this feature code to you order the Bureau of Standards, Metrology, and Inspection (BSMI) certification documents that are required when the storage system is shipped to Taiwan.

Table 48. Feature code for the BSMI certification documents (Taiwan)

Feature code	Description
0400	BSMI certification documents

Shipping weight reduction

Order the shipping weight reduction option to receive delivery of a storage system in multiple shipments.

If your site has delivery weight constraints, IBM offers a shipping weight reduction option that ensures the maximum shipping weight of the initial frame shipment does not exceed 909 kg (2000 lb). The frame weight is reduced by removing selected components, which are shipped separately.

The IBM service representative installs the components that were shipped separately during the storage system installation. This feature increases storage system installation time, so order it only if it is required.

Feature code for shipping weight reduction

Use this feature code to order the shipping-weight reduction option for your storage system.

This feature ensures that the maximum shipping weight of the base rack or expansion rack does not exceed 909 kg (2000 lb) each. Packaging adds 120 kg (265 lb).

Table 49. Feature code for shipping weight reduction

Feature code	Description	Models
0200	Shipping weight reduction	All

Chapter 5. Planning use of licensed functions

Licensed functions are the operating system and functions of the storage system. Required features and optional features are included.

IBM authorization for licensed functions is purchased as 239x machine function authorizations. However, the license functions are actually storage models. For example, the operating environment license (OEL) is listed as a 239x model LFA, OEL license (242x machine type). The 239x machine function authorization features are for billing purposes only.

Licensed function indicators

Each licensed function indicator feature that you order on a base frame enables that function at the system level.

After you receive and apply the feature activation codes for the licensed function indicators, the licensed functions are enabled for you to use. The license function indicators are also used for maintenance billing purposes.

Note: Retrieving feature activation codes is part of managing and activating your licenses. Before you can logically configure your storage system, you must first manage and activate your licenses.

Each licensed function indicator requires a corequisite 239x function authorization. Function authorization establishes the extent of IBM authorization for the licensed function before the feature activation code is provided by IBM. Each function authorization applies only to the specific storage system (by serial number) for which it was acquired. The function authorization cannot be transferred to another storage system (with a different serial number).

License scope refers to the following types of storage and types of servers with which the function can be used:

Fixed block (FB)

The function can be used only with data from Fibre Channel attached servers.

Count key data (CKD)

The function can be used only with data from FICON attached servers.

Both FB and CKD (ALL)

The function can be used with data from all attached servers.

You do not specify the license scope when you order function authorization. The function authorization establishes the extent of the IBM authorization (in terms of physical capacity), regardless of the storage type. However, if a licensed function has multiple license scope options, you must select a license scope when you initially retrieve the feature activation codes for your storage system. To select license scope, use the IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa).

When you use the DSFA website to change the license scope after a licensed function is activated, a new feature activation code is generated. When you install the new feature activation code into the storage system, the function is activated

and enforced by using the newly selected license scope. The increase in the license scope (changing FB or CKD to ALL) is a nondisruptive activity but takes effect at the next restart. A lateral change (changing FB to CKD or changing CKD to FB) or a reduction of the license scope (changing ALL to FB or CKD) is also a nondisruptive activity and takes effect at the next restart.

The following table lists feature codes for the licensed-function indicators and function authorization.

Note: Order these features only for base frames. Do not order these features for the expansion frames.

Table 50. Licensed function indicators for base frames

Licensed function	License scope	Feature code for licensed function indicator	Corequisite feature code for function-authorization
Operating environment	ALL	0700	7030 - 7065
Encrypted-drive activation	ALL	1750	
Encrypted-drive deactivation	ALL	1754	
FICON attachment	CKD	0703	7091
High Performance FICON	CKD	0709	7092
Database protection	FB, CKD, or ALL	0708	7080
FlashCopy	FB, CKD, or ALL	0720	7250 - 7260
Space Efficient FlashCopy (FlashCopy SE)	FB, CKD, or ALL	0730	7350 - 7360
Metro Mirror	FB, CKD, or ALL	0744	7500 - 7510
Multiple Target PPRC	FB, CKD, or ALL	0745	7025
Global Mirror	FB, CKD, or ALL	0746	7520 - 7530
Metro/Global Mirror	FB, CKD, or ALL	0742	7480 - 7490
z/OS Global Mirror	CKD	0760	7650 - 7660
z/OS Metro/Global Mirror Incremental Resync	CKD	0763	7680 - 7690
Parallel access volumes	CKD	0780	7820 - 7830
HyperPAV	CKD	0782	7899
Thin Provisioning	FB	0707	7071
IBM Easy Tier	FB, CKD, or ALL	0713	7083
Easy Tier Server	FB, CKD, or ALL	0715	7084
I/O Priority Manager	FB, CKD, or ALL	0784	7840 - 7850
z/OS Distributed Data Backup	CKD	0714	7094

License scope

Licensed functions are activated and enforced within a defined license scope.

License scope refers to the following types of storage and types of servers with which the function can be used:

Fixed block (FB)

The function can be used only with data from Fibre Channel attached servers.

Count key data (CKD)

The function can be used only with data from FICON attached servers.

Both FB and CKD (ALL)

The function can be used with data from all attached servers.

Some licensed functions have multiple license scope options, while other functions have only a single license scope. The following table provides the license scope options for each licensed function.

Table 51. License scope for each licensed function

Licensed function	License scope options
Operating environment	ALL
FICON attachment	CKD
High Performance FICON	CKD
Database protection	FB, CKD, or ALL
Point-in-time copy	FB, CKD, or ALL
Point-in-time copy	FB, CKD, or ALL
FlashCopy SE	FB, CKD, or ALL
FlashCopy SE	FB, CKD, or ALL
Remote mirror and copy	FB, CKD, or ALL
Global Mirror	FB, CKD, or ALL
Global Mirror	FB, CKD, or ALL
Metro Mirror	FB, CKD, or ALL
Metro Mirror	FB, CKD, or ALL
Multiple Target PPRC	FB, CKD, or ALL
Remote mirror for z/OS	CKD
Parallel access volumes	CKD
HyperPAV	CKD
Thin Provisioning	FB
IBM Easy Tier	FB, CKD, or All
Easy Tier Server	FB
I/O Priority Manager	FB, CKD
z/OS Distributed Data Backup	CKD
z/OS Global Mirror Incremental Resync	CKD

You do not specify the license scope when you order function authorization feature numbers. Feature numbers establish only the extent of the IBM authorization (in terms of physical capacity), regardless of the storage type. However, if a licensed function has multiple license scope options, you must select a license scope when you initially retrieve the feature activation codes for your storage system. This activity is performed by using the IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa).

Note: Retrieving feature activation codes is part of managing and activating your licenses. Before you can logically configure your storage system, you must first manage and activate your licenses.

When you use the DSFA website to change the license scope after a licensed function is activated, a new feature activation code is generated. When you install the new feature activation code into the storage system, the function is activated and enforced by using the newly selected license scope. The increase in the license scope (changing FB or CKD to ALL) is a nondisruptive activity but takes effect at the next restart. A lateral change (changing FB to CKD or changing CKD to FB) or a reduction of the license scope (changing ALL to FB or CKD) is also a nondisruptive activity and takes effect at the next restart.

Ordering licensed functions

After you decide which licensed functions to use with your storage system, you are ready to order the functions. Functions to include are operating environment license (OEL) features and optional licensed functions.

About this task

Licensed functions are purchased as function authorization features.

To order licensed functions, use the following general steps:

Procedure

1. Order the operating environment license (OEL) features that support the total physical capacity of your storage system.
2. Order optional licensed functions for your storage system.

Rules for ordering licensed functions

An operating environment license (OEL) is required for every base frame. All other licensed functions are optional.

For all licensed functions, you can combine feature codes to order the exact capacity that you need. For example, if you determine that you need 23 TB of point-in-time capacity, you can order two 7253 features (10 TB each) and three 7251 features (1 TB each).

Note: If you are activating features for any of the licensed functions, such as Copy Services, all features must have the same capacity, including the operating environment license feature.

When you calculate physical capacity, consider the capacity across the entire storage system, including the base frame and any expansion frames. To calculate the physical capacity, use The following figure to determine the total size of each regular drive feature in your storage system, and then add all the values.

Note: Standby CoD disk-drive features do not count toward the physical capacity.

Table 52. Total physical capacity for drive-set features

Drive sizes	Total physical capacity	Drives per feature
146 GB disk drives	2.336 TB	16

Table 52. Total physical capacity for drive-set features (continued)

Drive sizes	Total physical capacity	Drives per feature
200 GB flash drives	3.2 TB	16
300 GB disk drives	4.8 TB	16
400 GB flash drives	3.2 TB or 6.4 TB	8 or 16
400 GB flash cards	5.6 TB or 6.4 TB	14 or 16
600 GB disk drives	9.6 TB	16
800 GB flash drives	12.8 TB	16
900 GB disk drives	14.4 TB	16
1.2 TB disk drives	19.2 TB	16
1.6 TB disk drives	25.6 TB	16
3 TB disk drives	24 TB	8
4 TB disk drives	32 TB	8

Rules specific to 239x Model LFA, OEL license (machine type 242x)

The operating environment license (OEL) must cover the full physical capacity of your storage system, which includes the physical capacity of any expansion frame within the storage system. The license must cover both open systems data (fixed block data) and z Systems data (count key data). Standby CoD drives are not included in this calculation.

Note: Your storage system cannot be logically configured until you activate the OEL for it. Upon activation, drives can be logically configured up to the extent of the IBM OEL authorization level.

You can combine feature codes to order the exact capacity that you need. For example, if you determine that you need 25 TB of Metro Mirror capacity, you can order two 7503 features (10 TB each) and one 7502 feature (5 TB each).

As you add more drives to your storage system, you must increase the OEL authorization level for the storage system by purchasing more license features. Otherwise, you cannot logically configure the additional drives for use.

When you activate Standby CoD drives, you must also increase the OEL authorization to cover the activated Standby CoD capacity.

Rules specific to optional licensed functions

The following ordering rules apply when you order point-in-time licenses for Copy or remote mirror and copy licenses:

- If the function is used with only open systems data, a license is required for only the total physical capacity that is logically configured as fixed block (FB).
- If the function is used with only z Systems data, a license is required for only the total physical capacity that is logically configured as count key data (CKD).
- If the function is used for both open systems and z Systems data, a license is required for the total configured capacity.
- You must use Fibre Channel host adapters with remote mirror and copy functions. To see a current list of environments, configurations, networks, and

products that support remote mirror and copy functions, click **Interoperability Matrix** at the following IBM System Storage Interoperation Center (SSIC) website (www.ibm.com/systems/support/storage/config/ssic).

- You must purchase features for both the source (primary) and target (secondary) storage system.
- If you use the Metro/Global Mirror solution in your environment, the following rules apply:
 - Site A - You must have a Metro/Global Mirror license, and a Metro Mirror license.

Note: A Global Mirror Add-on license is required if you remove Site B and you want to resync between Site A and Site C.

- Site B - You must have a Metro/Global Mirror license, a Metro Mirror license, and a Global Mirror Add-on license.
- Site C - You must have a Metro/Global Mirror license, a Global Mirror license, and a point-in-time copy license.

A Metro/Global Mirror solution is available with the Metro/Global Mirror indicator feature numbers 74xx and 0742 and corresponding DS8000 series function authorization (2396-LFA MGM feature numbers 74xx).

- Site A - You must have a Metro/Global Mirror license, and a remote mirror and copy license.
- Site B - You must have a Metro/Global Mirror license, and a remote mirror and copy license.
- Site C - You must have a Metro/Global Mirror license, a remote mirror and copy license, and a point-in-time copy license.
- If you use Global Mirror, you must use the following rules:
 - A point-in-time copy function authorization (239x Model LFA, PTC license, 242x machine type) must be purchased for the secondary storage system.
 - If Global Mirror is to be used during failback on the secondary storage system, a point-in-time copy function authorization must also be purchased on the primary storage system.

The following ordering rule applies to remote mirror for z/OS licenses:

- A license is required for only the total physical capacity that is logically configured as count key data (CKD) volumes for use with z Systems host systems.
- When failback from the secondary storage system to the primary storage system is required, the remote mirror for z/OS function authorization (239x Model LFA, RMZ license, 242x machine type) must be purchased for both storage systems.

For parallel access volumes (PAV), a license is required for only the total physical capacity that is logically configured as count key data (CKD) volumes for use with z Systems host systems.

The following ordering rule applies to IBM HyperPAV:

- A license for IBM HyperPAV requires the purchase of PAV licensed features.

The initial enablement of any optional DS8000 licensed function is a concurrent activity (assuming that the appropriate level of microcode is installed on the machine for the specific function). The removal of a DS8000 licensed function to deactivate the function is a non-disruptive activity but takes effect at the next machine IML.

If you have an active optional function and no longer want to use it, you can deactivate the function in one of the following ways:

- Order an inactive or disabled license and replace the active license activation key with the new inactive license activation key at the IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa).
- Go to the DSFA website and change the assigned value from the current number of terabytes (TB) to 0 TB. This value, in effect, makes the feature inactive. If this change is made, you can go back to DSFA and reactivate the feature, up to the previously purchased level, without having to repurchase the feature.

Regardless of which method is used, the deactivation of a licensed function is a non-disruptive activity but takes effect at the next machine IML.

Note: Although you do not need to specify how the licenses are to be applied when you order them, you must allocate the licenses to the storage image when you obtain your license keys on the IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa).

Operating environment license (239x Model LFA, OEL license, 242x machine type)

The operating environment model and features establish the extent of IBM authorization for the use of the IBM DS operating environment.

For every storage system, you must order an operating environment license (OEL). This operating environment license support function is called the 239x Model LFA, OEL license on the 242x hardware machine type. The OEL licenses the operating environment and is based on the total physical capacity of the storage system (base frame plus any expansion frames). It authorizes you to use the model configuration at a specific capacity level. After the OEL is activated for the storage system, you can configure the storage system. Activating the OEL means that you obtained the feature activation key from the IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa) and entered it into the DS8000 Storage Management GUI.

Feature codes for the operating-environment license

Use these feature codes to order operating-environment licenses for your storage system. An operating-environment license is required per TB unit and per value unit for every storage system (including the base frame and all physically attached expansion frames). The extent of IBM authorization that is acquired through the function-authorization feature codes must cover the physical capacity and the value units for each drive that is installed in the storage system, excluding Standby CoD capacity.

The operating-environment license per value unit (model LFA 7050-7065) is required in addition to the operating-environment licenses per TB unit (feature codes 7030 - 7045). For each drive set, the corresponding number of value units must be purchased. The licensed machine code (LMC) does not allow the logical configuration of physical capacity beyond the extent of IBM authorization (except standby capacity on demand (CoD) capacity).

If the operating-environment license is not acquired and activated on the storage system, drives that are installed in the storage system cannot be logically

configured for use. After the operating environment is activated, drives can be logically configured up to the extent of IBM authorization.

As more drives are installed, the extent of IBM authorization must be increased by acquiring more function-authorization feature codes. Otherwise, the additional drives cannot be logically configured for use.

For a storage system with standby CoD disk drives (feature codes 5xx9), activation (logical configuration) of CoD disk drives can exceed the extent of IBM authorization for the operating environment, in which case an increased authorization level must be acquired.

Note: The following activities are non-disruptive and take effect at the next machine IML.

- A lateral change, in which the license scope is changed from fixed block (FB) to count key data (CKD) or from CKD to FB
- A reduction in the license scope, in which license scope is changed from all physical capacity (ALL) to only FB or only CKD capacity
- A deactivation of an activated licensed function

The following table lists the feature codes for operating-environment licenses.

Table 53. Feature codes for operating-environment licenses

Operating environment	Feature code for licensed function indicator	Corequisite feature code for function-authorization
Inactive	0700	7030
1 TB unit	0700	7031
5 TB unit	0700	7032
10 TB unit	0700	7033
25 TB unit	0700	7034
50 TB unit	0700	7035
100 TB unit	0700	7040
200 TB unit	0700	7045
Inactive	0700	7050
1 value unit	0700	7051
5 value unit	0700	7052
10 value unit	0700	7053
25 value unit	0700	7054
50 value unit	0700	7055
100 value unit	0700	7060
200 value unit	0700	7065

Parallel access volumes (239x Model LFA, PAV license; 242x machine type)

The parallel access volumes (PAV) features establish the extent of IBM authorization for the use of the parallel access volumes licensed function.

Feature codes for parallel access volume licensed function

Use these feature codes to order parallel access volume (PAV) licensed function for the base frame (model 961).

A license is required for the total physical capacity in the entire storage system (base frame and all attached expansion frames) that is configured as count key data (CKD). The total authorization level must be greater than or equal to the total physical capacity of the storage system.

When you order the PAV function, you must specify the feature code that represents the total physical capacity that is configured as CKD. You can combine feature codes to order the exact capacity that you need. For example, if you need a function authorization for 35 TB, you would order one 7824 feature and one 7822 feature.

Note: If you currently have an active PAV feature, and you replace it with an inactive feature, but later want to use the feature again, adhere to the requirements for deleting an active license.

The following table lists the feature codes for the PAV licensed function.

Table 54. Feature codes for the parallel access volume

License function indicator	Feature code for license function indicator	Corequisite feature code for function-authorization
Inactive	0780	7820
1 TB unit	0780	7821
5 TB unit	0780	7822
10 TB unit	0780	7823
25 TB unit	0780	7824
50 TB unit	0780	7825
100 TB unit	0780	7830

IBM HyperPAV (242x Model PAV and 239x Model LFA, PAV license)

You can add the optional IBM HyperPAV feature to any licensed parallel access volume (PAV) feature.

IBM HyperPAV can be enabled only if PAV is enabled on the storage system. The IBM HyperPAV feature is available for a single charge (flat fee) regardless of the extent of IBM authorization that you have for the corresponding PAV feature.

Feature code for IBM HyperPAV licensed function

Use this feature code to order the IBM HyperPAV function for an existing or new parallel access volumes (PAV) on your storage system.

Table 55. Feature code for IBM HyperPAV licensed function

License function indicator	Feature code for license function indicator	Corequisite feature code for function-authorization (239x-LFA)
IBM HyperPAV	0782	7899

IBM Easy Tier

Support for IBM Easy Tier is available with the IBM Easy Tier licensed feature.

The Easy Tier license feature enables the following modes:

- Easy Tier: automatic mode
- Easy Tier: manual mode

The license feature enables the following functions for the storage type:

- Easy Tier application
- Easy Tier heat map transfer
- The capability to migrate volumes for logical volumes
- The reconfigure extent pool function of the extent pool
- The dynamic extent relocation with an Easy Tier managed extent pool

The Easy Tier licensed feature key contains a storage-type indication that determines the type of storage for which the key is applicable. It also contains an allowed capacity value. This value refers to the total amount of physical capacity that is configured into any real rank of the specified storage types in the storage system. The allowed capacity is required to be set to either 0 or to the maximum value, indicating whether the licensed feature is on or off.

To validate an Easy Tier licensed feature key, the allowed capacity indication must meet all of the following criteria:

- The specified storage type must be fixed block (FB) or count key data (CKD).
- The specified capacity must be either zero or the maximum capacity value.

When an Easy Tier licensed feature key is installed, if the Easy Tier functions are not enabled and the license feature key has a capacity greater than 0 bytes, then the storage system enables the Easy Tier functions.

If the licensed feature key that is disabled, while the Easy Tier functions are enabled, the disabled licensed feature key is accepted and the Easy Tier functions are disabled immediately. Any extent migrations that were in progress during disablement are either nullified or completed. Any extent migrations that are queued later are stopped and any requests to initiate a volume migration or an extent pool reconfiguration are rejected.

Feature codes for IBM Easy Tier licensed function

Use these feature codes to order the IBM Easy Tier licensed function. The license enables the use of Easy Tier for your storage system.

Table 56. Feature codes for IBM Easy Tier licensed function

License function	Feature code for licensed function indicator	Corequisite feature code for function-authorization
IBM Easy Tier	0713	7083

Feature codes for IBM Easy Tier Server licensed function

Use this feature code to order IBM Easy Tier Server function for your storage system. The license enables the use of Easy Tier Server for your storage system.

To validate an Easy Tier Server licensed function, the storage type must be fixed block (FB).

Table 57. Feature codes for Easy Tier Server licensed function

License function	Feature code for licensed function indicator	Corequisite feature code for function-authorization
Easy Tier Server	0715	7084

Point-in-time copy function (239x Model LFA, PTC license) and FlashCopy SE Model SE function (239x Model LFA, SE license)

The point-in-time copy licensed function model and features establish the extent of IBM authorization for the use of the point-in-time copy licensed function on your storage system.

The IBM FlashCopy function is a point-in-time licensed function.

Feature codes for FlashCopy licensed function

Use these feature codes to order the FlashCopy licensed function for the base frame (model 961). The FlashCopy license enables the use of FlashCopy for your storage system.

Note: If you are activating features for any of the licensed functions, such as Copy Services, all features must have the same capacity, including the operating environment license feature.

When you order the FlashCopy licensed function, you must specify the feature codes that represents the physical capacity that you want to authorize for FlashCopy. You can combine feature codes to order the exact capacity that you need. For example, if you need 23 TB of capacity, you would order two 7253 features and three 7251 features.

Note: If you have an active FlashCopy feature and replace it with an inactive feature, but later want to use the feature again, adhere to the requirements for deleting an active license.

The following table lists the feature codes for the FlashCopy licensed function.

Table 58. Feature codes for FlashCopy licensed function

Licensed function	Feature code for licensed function indicator	Corequisite feature code for function-authorization
PTC - Inactive	0720	7250
PTC - 1 TB unit	0720	7251
PTC - 5 TB unit	0720	7252
PTC - 10 TB unit	0720	7253
PTC - 25 TB unit	0720	7254
PTC - 50 TB unit	0720	7255
PTC - 100 TB unit	0720	7260

Feature codes for Space Efficient FlashCopy licensed function

Use these feature codes to order the Space Efficient FlashCopy (FlashCopy SE) licensed function for the base frame (model 961). The FlashCopy SE license enables the use of the FlashCopy SE for your storage system.

Note: If you are activating features for any of the licensed functions, such as Copy Services, all features must have the same capacity, including the operating environment license feature.

When you order the FlashCopy SE licensed function, you must specify the feature codes that represents the physical capacity that you want to authorize for FlashCopy SE. You can combine feature codes to order the exact capacity that you need. For example, if you need 23 TB of capacity, you would order two 7353 features and three 7351 features.

Note: If you have an active FlashCopy SE feature and replace it with an inactive feature, but later want to use the feature again, adhere to the requirements for deleting an active license.

The following table lists feature codes for the FlashCopy SE licensed function.

Table 59. Feature codes for FlashCopy licensed function

Licensed function	Feature code for licensed function indicator	Corequisite feature code for function-authorization
Inactive	0730	7350
1 TB unit	0730	7351
5 TB unit	0730	7352
10 TB unit	0730	7353
25 TB unit	0730	7354
50 TB unit	0730	7355
100 TB unit	0730	7360

Remote mirror and copy functions (242x Model RMC and 239x Model LFA)

The remote mirror and copy licensed function model and features establish the extent of IBM authorization for the use of the remote mirror and copy licensed functions on your storage system.

The following functions are remote mirror and copy licensed functions:

- Metro Mirror
- Global Mirror
- Global Copy
- Metro/Global Mirror
- Multiple Target PPRC

Feature codes for remote mirror and copy

Use these feature codes to order the remote mirror and copy licensed function for the base frame (model 961).

The remote mirror and copy license feature codes enable the use of the following remote mirror and copy licensed functions:

- IBM Metro Mirror (MM)
- IBM Global Mirror (GM)
- IBM Metro Global Mirror (MGM)

Note: If you are activating features for any of the licensed functions, such as Copy Services, all features must have the same capacity, including the operating environment license feature.

When you order remote mirror and copy licensed functions, you must specify the feature code that represents the physical capacity to authorize for the function. You can combine feature codes to order the exact capacity that you need. For example, if you determine that you need a function authorization for 35 TB of MM capacity, you would order one 7504 feature and one 7503 feature.

Note: If you have an active remote mirror and copy feature and replace it with an inactive feature, but later want to use the feature again, adhere to the requirements for deleting an active license.

The following table lists the feature codes for the remote mirror and copy licensed function.

Table 60. Feature codes for remote mirror and copy licensed function

Licensed function	Feature code for licensed function indicator	Corequisite feature code for function-authorization
MGM - inactive	0742	7480
MGM - 1 TB unit	0742	7481
MGM - 5 TB unit	0742	7482
MGM - 10 TB unit	0742	7483
MGM - 25 TB unit	0742	7484
MGM - 50 TB unit	0742	7485
MGM - 100 TB unit	0742	7490
MM - Inactive	0744	7500
MM - 1 TB unit	0744	7501
MM - 5 TB unit	0744	7502
MM - 10 TB unit	0744	7503
MM - 25 TB unit	0744	7504
MM - 50 TB unit	0744	7505
MM - 100 TB unit	0744	7510
Multiple Target PPRC	0745	7025
GM - Inactive	0746	7520
GM - 1 TB unit	0746	7521
GM - 5 TB unit	0746	7522
GM - 10 TB unit	0746	7523
GM - 25 TB unit	0746	7524
GM - 50 TB unit	0746	7525

Table 60. Feature codes for remote mirror and copy licensed function (continued)

Licensed function	Feature code for licensed function indicator	Corequisite feature code for function-authorization
GM - 100 TB unit	0746	7530

Feature codes for I/O Priority Manager

Use these feature codes to order I/O Priority Manager for the DS8870 base model.

Note: If you are activating features for any of the licensed functions, such as Copy Services, all features must have the same capacity, including the operating environment license feature.

Select the feature code that represents the physical capacity to authorize for the function. You can combine feature codes to order the exact capacity that you need. For example, if you need a function authorization for 35 TB of I/O Priority Manager capacity, you would order one 7843 feature and one 7844 feature.

The following table lists the feature codes for the I/O Priority Manager functions.

Table 61. Feature codes for I/O Priority Manager

License function	Feature code for licensed function indicator	Corequisite feature code for function-authorization (239x-LFA)
Inactive	0784	7840
1 TB indicator	0784	7841
5 TB indicator	0784	7842
10 TB indicator	0784	7843
25 TB indicator	0784	7844
50 TB indicator	0784	7845
100 TB indicator	0784	7850

z/OS licensed features

This section describes z/OS licensed features that are supported on the storage system.

Remote mirror for z/OS (242x Model RMZ and 239x Model LFA, RMZ license)

The remote mirror for z/OS licensed function model and features establish the extent of IBM authorization for the use of the z/OS remote mirroring licensed function on your storage system.

The IBM z/OS Global Mirror function is a z/OS remote mirroring licensed function.

Feature codes for z/OS Global Mirror licensed function

Use these feature codes to order the z/OS Global Mirror (RMZ) licensed function for the base frame (model 961).

Note: If you are activating features for any of the licensed functions, such as Copy Services, all features must have the same capacity, including the operating environment license feature.

When you order the RMZ function, you must specify the feature code that represents the physical capacity you want to authorize for this function. You can combine feature codes to order the exact capacity that you need. For example, if you determine that you need 30 TB of capacity, you would order one 7654 feature and one 7652 feature.

Note: If you have an active RMZ feature and replace it with an inactive feature, but later want to use the feature again, adhere to the requirements for deleting an active license.

The following table lists the feature codes for remote mirror for z/OS functions.

Table 62. Feature codes for the z/OS Global Mirror licensed function

Licensed function	Feature code for licensed function indicator	Corequisite feature code for function-authorization
Inactive	0760	7650
1 TB unit	0760	7651
5 TB unit	0760	7652
10 TB unit	0760	7653
25 TB unit	0760	7654
50 TB unit	0760	7655
100 TB unit	0760	7660

Feature codes for z/OS Metro/Global Mirror Incremental Resync licensed function

Use these feature codes to order the z/OS Metro/Global Mirror Incremental Resync (z/Global Mirror Resync) licensed function for the base frame (model 961).

Note: If you are activating features for any of the licensed functions, such as Copy Services, all features must have the same capacity, including the operating environment license feature.

When you order the z/Global Mirror Resync function, you must specify the feature code that represents the physical capacity you are authorizing for the function. You can combine feature codes to order the exact capacity that you need. For example, if you determine that you need 30 TB of capacity, you would order one 7684 feature and one 7682 feature.

Note: If you have an active z/Global Mirror Resync feature and replace it with an inactive feature, but later want to use the feature again, adhere to the requirements for deleting an active license.

The following table lists the feature codes for z/Global Mirror Resync) licensed function.

Table 63. Feature codes for the z/OS Metro/Global Mirror Incremental Resync licensed function

Licensed function indicator	Feature code for licensed function indicator	Corequisite feature code for function-authorization
Inactive	0763	7680
1 TB unit	0763	7681
5 TB unit	0763	7682
10 TB unit	0763	7683
25 TB unit	0763	7684
50 TB unit	0763	7685
100 TB unit	0763	7690

z/OS Distributed Data Backup

z/OS Distributed Data Backup (zDDB) is an optional licensed feature on the base frame that allows hosts, which are attached through a FICON or ESCON interface, to access data on fixed block (FB) volumes through a device address on FICON or ESCON interfaces.

If the zDDB licensed feature key is installed and enabled and a volume group type specifies either FICON or ESCON interfaces, this volume group has implicit access to all FB logical volumes that are configured in addition to all CKD volumes specified in the volume group. Then, with appropriate software, a z/OS host can complete backup and restore functions for FB logical volumes that are configured on a storage system image for open systems hosts.

The hierarchical storage management (DFSMSHsm) function of the z/OS operating system can manage data that is backed up. If the zDDB licensed feature key is not installed or enabled, during a storage system power-on sequence, the logical volumes, and the LSSs that are associated with this licensed feature are offline to any FICON or ESCON interfaces. If a zDDB licensed feature key is disabled when it was previously enabled, the logical volumes and the LSSs that are associated with the licensed features remain online to any FICON or ESCON interfaces until the next power off sequence, but any I/O that is issued to these logical volumes is rejected.

The key data in a zDDB LIC feature key contains an allowed capacity value. This value refers to the total amount of physical capacity that is configured into any FB real rank on the storage system image. The allowed capacity is required to be set to either 0 or to the maximum value, which the LIC feature is on or off.

To validate a zDDB feature key, the new license feature key must meet the following criteria:

- The specified storage type must be FB storage only
- The specified capacity is either zero or the maximum capacity value

When a zDDB licensed feature key is installed and the licensed feature key has a capacity greater than 0 bytes, then the storage system enables the zDDB function and notifies any attached hosts through the appropriate interface. If a zDDB licensed feature key that is disabled is installed while the zDDB facility is enabled, the licensed feature key is accepted, but the zDDB facility is concurrently disabled.

While the feature is disabled, logical volumes that are associated with this feature are either offline to FICON or ESCON hosts or I/O that is issued to logical volumes associated with the feature are rejected on FICON or ESCON interfaces.

Feature codes for z/OS Distributed Data Backup licensed function

Use this feature code to order the z/OS Distributed Data Backup licensed function for your storage system. You can use z/OS Distributed Data Backup to back up data for open systems from distributed server platforms through a z Systems host.

Table 64. Feature code for the z/OS Distributed Data Backup licensed function

License function indicator	Feature code for license function indicator	Corequisite feature code for function-authorization (239x-LFA)
z/OS Distributed Data Backup Indicator	0714	7094

Thin provisioning licensed feature key

To use the thin provisioning facility with extent space-efficient logical volumes, you must have the thin provisioning licensed feature key.

The thin provisioning licensed feature enables the following functions for the storage type that is indicated by the licensed feature key:

- The creation of extent space-efficient logical volumes
- The creation of virtual ranks

The thin provisioning licensed feature key data contains the following information:

- A storage-type indication that determines the type of storage that is applicable for the licensed feature key
- A capacity value that refers to the total amount of physical capacity that is configured into any real rank for the storage types in the storage system.

For the thin provisioning licensed feature key to be valid, the capacity value must meet all of the following criteria:

- The specified storage type must be fixed block (FB)
- The specified capacity must be either zero or the maximum capacity value

The support of FB thin provisioning depends on the model. If thin provisioning is not supported on all supported storage types, the licensed feature key must still indicate both storage types. A request to create either an extent space-efficient logical volume or a virtual rank in a storage type that is not supported is rejected. Subsequent changes to storage types that are supported do not require a new licensed feature key to be installed. Support of thin provisioning is only indicated to the host types that access storage types that provide support for thin provisioning.

When a thin provisioning licensed feature key is installed, if the thin provisioning facility is not currently enabled and the licensed feature key has a capacity greater than 0 bytes, the storage facility image enables the thin provisioning facility and notifies any attached hosts through the appropriate interface protocol to single this condition. If there are one or more extent space-efficient logical volumes that are configured on the storage system and the installed licensed feature key is disabled while the thin provisioning facility is enabled, the licensed feature key is rejected. Otherwise, the following results occur:

- The disablement licensed feature key is accepted
- The thin provisioning facility is immediately disabled
- Any capabilities that were previously enabled for the thin provisioning facility are disabled

Note: If a capability is enabled by some other licensed feature key, it remains enabled.

- All attached hosts are notified through the appropriate interface protocol that the thin provisioning facility is disabled

If configuration-based enforcement is in effect while the thin provisioning facility is enabled by a licensed feature key, the configuration of more real ranks of the specified storage types is suppressed on the storage system if the resulting physical capacity of the specified storage types can exceed the amount that is allowed by the licensed feature key.

Note: Thin provisioning functions are not supported on System z/OS volumes.

Extent Space Efficient (ESE) capacity controls for thin provisioning

Use of thin provisioning can affect the amount of storage capacity you choose to order. ESE capacity controls allow you to allocate storage appropriately.

With the mixture of thin-provisioned (ESE) and fully-provisioned (non-ESE) volumes in an extent pool, a method is needed to dedicate some of the extent-pool storage capacity for ESE user data usage as well as limit the ESE user data usage within the extent pool. Also needed is the ability to detect when the available storage space within the extent pool for ESE volumes is running out of space.

ESE capacity controls provide extent pool attributes to limit the maximum extent pool storage available for ESE user data usage, and to guarantee a proportion of the extent pool storage to be available for ESE user data usage.

Associated with the ESE capacity controls is an SNMP trap that notifies you when the ESE extent usage in the pool exceeds an ESE extent threshold set by you, as well as when the extent pool is out of storage available for ESE user data usage.

ESE capacity controls include the following attributes:

ESE Extent Threshold

The percentage that is compared to the actual percentage of storage capacity available for ESE customer extent allocation when determining the extent pool ESE extent status.

ESE Extent Status

One of the three following values:

- 0: the percent of the available ESE capacity is greater than the ESE extent threshold
- 1: the percent of the available ESE capacity is greater than zero but less than or equal to the ESE extent threshold
- 10: the percent of the available ESE capacity is zero

Note: When the size of the extent pool remains fixed or is only increased, the allocatable physical capacity remains greater than or equal to the allocated physical

capacity. However, a reduction in the size of the extent pool can cause the *allocatable* physical capacity to become less than the *allocated* physical capacity in some cases.

For example, if the user requests that one of the ranks in an extent pool be depopulated, the data on that rank are moved to the remaining ranks in the pool causing the rank to become not allocated and removed from the pool. The user is advised to inspect the limits and threshold on the extent pool following any changes to the size of the extent pool to ensure that the specified values are still consistent with the user's intentions.

Feature codes for thin provisioning licensed function

Use this feature code to order the thin provisioning function for your storage system.

Table 65. Feature code for the thin provisioning license function

License function	Feature code for licensed function indicator	Corequisite feature code for function-authorization
Thin provisioning indicator	0707	7071

Chapter 6. Meeting delivery and installation requirements

You must ensure that you properly plan for the delivery and installation of your storage system.

This information provides the following planning information for the delivery and installation of your storage system:

- Planning for delivery of your storage system
- Planning the physical installation site
- Planning for power requirements
- Planning for network and communication requirements

For more information about the equipment and documents that IBM includes with storage system shipments, see Appendix C, “IBM DS8000 equipment and documents,” on page 211.

Delivery requirements

Before you receive your storage system shipment, ensure that the final installation site meets all delivery requirements.

Attention: Customers must prepare their environments to accept the storage system based on this planning information, with assistance from an IBM Advanced Technical Services (ATS) representative or an IBM service representative. The final installation site within the computer room must be prepared *before* the equipment is delivered. If the site cannot be prepared before the delivery time, customers must make arrangements to have the professional movers return to finish the transportation later. Only professional movers can transport the equipment. The IBM service representative can minimally reposition the frame at the installation site, as needed to complete required service actions. Customers are also responsible for using professional movers in the case of equipment relocation or disposal.

Acclimation

Server and storage equipment must be acclimated to the surrounding environment to prevent condensation.

When server and storage equipment is shipped in a climate where the outside temperature is below the dew point of an indoor location, water condensation might form on the cooler surfaces inside the equipment when brought into a warmer indoor environment. If condensation occurs, sufficient time must be allowed for the equipment to reach equilibrium with the warmer indoor temperature before you power on the storage system for installation. Leave the storage system in the shipping bag for a minimum of 24 to 48 hours to let it acclimate to the indoor environment.

Shipment weights and dimensions

To help you plan for the delivery of your storage system, you must ensure that your loading dock and receiving area can support the weight and dimensions of the packaged storage system shipments.

You receive at least two, and up to three, shipping containers for each frame that you order. You always receive the following items:

- A container with the storage system frame. In the People's Republic of China (including Hong Kong S.A.R. of China), India, and Brazil, this container is a wooden crate. In all other countries, this container is a pallet that is covered by a corrugated fiberboard (cardboard) cover.
- A container with the remaining components, such as power cords, CDs, and other ordered features or peripheral devices for your storage system

If ordered, you also receive a container with the external management consoles (MCs).

Table 66 shows the final packaged dimensions and maximum packaged weight of the storage system frame shipments.

To calculate the weight of your total shipment, add the weight of each frame container and the weight of one ship group container for each frame. If you ordered any external management consoles, add the weight of those containers, as well.

Table 66. Packaged dimensions and weight for storage system frames (all countries)

Container	Packaged dimensions	Maximum packaged weight
DS8870 All-Flash base frame (model 961)	Height 207.5 cm (81.7 in.) ¹ Width 101.5 cm (40 in.) Depth 137.5 cm (54.2 in.)	1395 kg (3075 lb)
DS8870 base frame (model 961)	Height 207.5 cm (81.7 in.) ¹ Width 101.5 cm (40 in.) Depth 137.5 cm (54.2 in.)	1451kg (3200 lb)
DS8870 expansion frame (model 96E)	Height 207.5 cm (81.7 in.) Width 101.5 cm (40 in.) Depth 137.5 cm (54.2 in.)	1279 kg (2820 lb)
External management console container (when ordered)	Height 69.0 cm (27.2 in.) Width 80.0 cm (31.5 in.) Depth 120.0 cm (47.3 in.)	75 kg (165 lb)
Note:		
1. With an overhead cabling option (top exit bracket, feature code 1400) installed on the base frame, an extra 10.16 cm (4 inches) are added to the standard, packaged height of the base frame. The overhead cabling option increases the total height of the frame to 87.6 cm (217.5 inches).		

Receiving delivery

The shipping carrier is responsible for delivering and unloading the storage system as close to its final destination as possible. You must ensure that your loading ramp and your receiving area can accommodate your storage system shipment.

About this task

Use the following steps to ensure that your receiving area and loading ramp can safely accommodate the delivery of your storage system:

Procedure

1. Find out the packaged weight and dimensions of the shipping containers in your shipment.
2. Ensure that your loading dock, receiving area, and elevators can safely support the packaged weight and dimensions of the shipping containers.

Note: You can order a weight-reduced shipment when a configured storage system exceeds the weight capability of the receiving area at your site.

3. To compensate for the weight of the storage system shipment, ensure that the loading ramp at your site does not exceed an angle of 10° . (See Figure 13.)




Figure 13. Maximum tilt for a packed frame is 10°

Installation site requirements

You must ensure that the location where you plan to install your storage system meets all requirements.

Planning for floor and space requirements

Ensure that the location where you plan to install your storage system meets space and floor requirements. Decide whether your storage system is to be installed on a raised or nonraised floor.

About this task

When you are planning the location of your storage system, you must answer the following questions that relate to floor types, floor loads, and space:

- What type of floor does the installation site have? The storage system can be installed on a raised or nonraised floor.
- If the installation site has a raised floor, does the floor require preparation (such as cutting out tiles) to accommodate cable entry into the system?
- Does the floor of the installation site meet floor-load requirements?
- Can the installation site accommodate the amount of space that is required by the storage system, and does the space meet the following criteria?
 - Weight distribution area that is needed to meet floor load requirements
 - Service clearance requirements

- Does the installation site require overhead cable management for host fiber and power cables?

Use the following steps to ensure that your planned installation site meets space and floor load requirements:

Procedure

1. Identify the base frame and expansion frames that are included in your storage system. If your storage system uses external management consoles, include the racks that contain the external management consoles.
2. Decide whether to install the storage system on a raised or nonraised floor.
 - a. If the location has a raised floor, plan where the floor tiles must be cut to accommodate the cables.
 - b. If the location has a nonraised floor, resolve any safety problems, and any special equipment considerations, caused by the location of cable exits and routing.
3. Determine whether the floor of the installation site meets the floor load requirements for your storage system.
4. Calculate the amount of space to be used by your storage system.
 - a. Identify the total amount of space that is needed for your storage system by using the dimensions of the frames and the weight distribution areas that are calculated in step 3.
 - b. Ensure that the area around each frame and each storage system meets the service clearance requirements.

Note: Any expansion frames in the storage system must be attached to the base frame on the right side as you face the front of the storage system.

Installing on raised or nonraised floors

You can install your storage system on a raised or nonraised floor. Raised floors can provide better cooling than nonraised floors.

Raised floor considerations

Installing your storage system on a raised floor provides the following benefits:

- Improves operational efficiency and allows greater flexibility in the arrangement of equipment.
- Increases air circulation for better cooling.
- Protects the interconnecting cables and power receptacles.
- Prevents tripping hazards because cables can be routed underneath the raised floor.

When you install a raised floor, consider the following factors:

- The raised floor must be constructed of fire-resistant or noncombustible material.
- The raised-floor height must be at least 30.5 cm (12 in.). For processors with multiple channels, a minimum raised-floor height of 45.7 cm (18 in.) is required. Clearance must be adequate to accommodate interconnecting cables, Fibre Channel cable raceways, power distribution, and any piping that is present under the floor. Floors with greater raised-floor heights allow for better equipment cooling.
- Fully configured, two-frame storage systems can weigh in excess of 2370 kg (5220 lbs). You must ensure that the raised floor on which the storage system is to be installed is able to support this weight. Contact the floor-tile

manufacturer and a structural engineer to verify that the raised floor is safe to support the concentrated loads equal to one third of the total weight of one frame. Under certain circumstances such as relocation, it is possible that the concentrated loads can be as high as one half of the total weight of one frame per caster. When you install two adjacent frames, it is possible that two casters induce a total load as high as one third of the total weight of two adjacent frames.

- Depending on the type of floor tile, more supports (pedestals) might be necessary to maintain the structural integrity of an uncut panel or to restore the integrity of a floor tile that is cut for cable entry or air supply. Contact the floor-tile manufacturer and a structural engineer to ensure that the floor tiles and pedestals can sustain the concentrated loads.
- Pedestals must be firmly attached to the structural (concrete) floor by using an adhesive.
- Seal raised-floor cable openings to prevent chilled air that is not used to directly cool the equipment from escaping.
- Use noncombustible protective molding to eliminate sharp edges on all floor cutouts, to prevent damage to cables and hoses, and to prevent casters from rolling into the floor cutout.
- Avoid the exposure of metal or highly conductive material at ground potential to the walking surface when a metallic raised floor structure is used. Such exposure is considered an electrical safety hazard.
- Concrete subfloors require treatment to prevent the release of dust.
- The use of a protective covering (such as plywood, tempered masonite, or plyron) is required to prevent damage to floor tiles, carpeting, and tiles while equipment is being moved to or is relocated within the installation site. When the equipment is moved, the dynamic load on the casters is greater than when the equipment is stationary.

Nonraised floor considerations

For environments with nonraised floors, an optional overhead cabling feature is available.

Follow the special considerations and installation guidelines as described in the topics about overhead cable management.

When you install a storage system on a non-raised floor, consider the following factors:

- The use of a protective covering (such as plywood, tempered masonite, or plyron) is required to prevent damage to floor and carpeting while equipment is being moved to or is relocated within the installation site.
- Concrete floors require treatment to prevent the release of dust.

Overhead cable management (top-exit bracket)

Overhead cable management (top-exit bracket) is an optional feature that includes top-exit bracket for managing your Fibre cables. This feature is an alternative to the standard, floor-cable exit.

Using overhead cabling provides many of the cooling and safety benefits that are provided by raised flooring in a nonraised floor environment. Unlike raised-floor cabling, the installation planning, cable length, and the storage-system location in relation to the cable entry point are critical to the successful installation of the top-exit bracket.

Figure 14 on page 145 illustrates the location of the cabling for the top-exit bracket for fiber cable feature. When you order the overhead-cable management feature, the feature includes clamping hardware, internal cable routing brackets for rack 1 or rack 2, and two top-exit mainline power cords for each rack. The following notes provide more information about the color-coded cable routing and components in Figure 14 on page 145.

1 Customer Fibre Channel host cables. The Fibre Channel host cables, which are shown in red, are routed from the top of the rack down to I/O enclosure host adapters.

2 Network Ethernet cable, power sequence cables, and customer analog phone line (if used). The network Ethernet cable, in blue, is routed from the top of rack to the rear rack connector. The rack connector has an internal cable to the management console. The power sequence cables and private network Ethernet cables (one gray and one black) for partner storage system or external management console (if installed) are also located here.

3 Mainline power cords. Two top-exit mainline power cords for each rack, which is shown in green, are routed here.

Notes:

- An IBM service representative installs and tests the power sources. The customer is required to provide power outlets (for connecting power cords) within the specified distance.
- Fibre Channel host cables are internally routed and connected by either the customer or by an IBM service representative.
- All remaining cables are internally routed and connected by an IBM service representative.




Figure 14. Top exit feature installed (cable routing and top exit locations)

Feature codes for overhead cable management (top-exit bracket):

Use this feature code to order cable management for overhead cabling (top exit bracket) for your storage system.

Note: In addition to the top-exit bracket and top-exit power cords, one IBM approved ladder (feature code 1101) must also be purchased for a site where the top-exit bracket for fiber cable feature is used. The IBM approved ladder is used to ensure safe access when your storage system is serviced with a top-exit bracket feature installed.

Table 67. Feature codes for the overhead cable (top-exit bracket)

Feature Code	Description
1400	Top-exit bracket for fiber cable

Overhead cabling installation and safety requirements:

Ensure that installation and safety requirements are met before your storage system is installed.

If the cables are too long, there is not enough room inside of the rack to handle the extra length and excess cable might interfere with the service process, preventing concurrent repair. Consider the following specifications and limitations before you order this feature:

- In contrast to the raised-floor power cords, which have a length from the tailgate to the connector of about 4.9 m (16 ft), the length of the top exit power cords are only 1.8 m (6 ft) from the top of the storage system.
- IBM Corporate Safety restricts the servicing of your overhead equipment to a maximum of 3 m (10 ft) from the floor. Therefore, your power source must not exceed 3 m (10 ft) from the floor and must be within 1.5 m (5 ft) of the top of the power cord exit gate. Servicing any overhead equipment higher than 3 m (10 ft) requires a special bid contract. Contact your IBM service representatives for more information on special bids.
- To meet safety regulations in servicing your overhead equipment, you must purchase a minimum of one feature code 1101 for your top exit bracket feature per site. This feature code provides a safety-approved 8-foot platform ladder, which is required to service feature codes 1072, 1073, 1083, 1084, and 1400. This ladder provides IBM service representatives the ability to perform power safety checks and other service activities on the top of your storage system. Without this approved ladder, IBM service representatives are not able to install or service a storage system with the top-cable exit features.
- To assist you with the top-exit host cable routing, feature code 1400 provides a cable channel bracket that mounts directly below the topside of the tailgate and its opening. Cables can be easily slid into the slots on its channels. The cable bracket directs the cables behind the rack ID card and towards the rear, where the cables drop vertically into a second channel, which mounts on the left-side wall (when viewing the storage system from the rear). There are openings in the vertical channel for cables to exit toward the I/O drawers.

Accommodating cables

You must ensure that the location and dimensions of the cable cutouts for each frame in the storage system can be accommodated by the installation location. An overhead-cable management option (top-exit bracket) is available for DS8870 for environments that have special planning and safety requirements.

Use the following steps to ensure that you prepare for cabling for each storage system:

1. Based on your planned storage system layout, ensure that you can accommodate the locations of the cables that exit each frame. See the following figure for the cable cutouts for the DS8870.




Figure 15. Cable cutouts for DS8870

2. If you install the storage system on a raised floor, use the following measurements when you cut the floor tile for cabling:
 - Width: 45.7 cm (18.0 in.)
 - Depth: 16 cm (6.3 in.)

Note: If both frames 1 and 2 use an overhead-cable management (top-exit bracket) feature for the power cords and communication cables, the PCIe and SPCN cables can be routed under the frame, on top of the raised floor. This is the same routing that is used for nonraised floor installations. There is room under the frame to coil extra cable length and prevent the need for custom floor tile cutouts. Also, frames 3 and 4 do not need floor tile cutouts when the top-exit bracket feature is installed, as only routing for the power cords is needed.

Nonraised floors with overhead cable management

Raised floors are recommended to provide better support for the cabling that is needed by the storage systems, and to ensure that you have efficient cooling for your storage system. However, for the base frame, an option is available for overhead cabling by using the top exit bracket feature, which provides many benefits for nonraised floor installations. Unlike raised-floor cabling, the installation planning, cable length, and the storage-system location in relation to the cable entry point are critical to the successful installation of a top-exit bracket feature. Measurements for this feature are given in the following figure. You can find

critical safety, service, and installation considerations for this feature in the topic that discusses overhead-cable management.

The following figure illustrates the location of these components:

- 1** Top exit bracket for fiber cables
- 2** Top exit cable channel bracket




FIG01681

Figure 16. Measurements for DS8870 with top exit bracket feature present

Physical footprint

The physical footprint dimensions, caster locations, and cable openings of the storage system help you plan your installation site.

The following figure shows the overall physical footprint of a storage system. The following dimensions are labeled on figure:

- 1** Front cover width
- 2** Front service clearance
- 3** Back cover widths

- 4** Back service clearance
- 5** Clearance to allow front cover to open
- 6** Distance between casters
- 7** Depth of frame without covers
- 8** Depth of frame with covers
- 9** Minimum dimension between casters and outside edges of frames
- 10** Distance from the edge to the front of the open cover




Figure 17. Physical footprint. Dimensions are in centimeters (inches).

Meeting floor load requirements

It is important for your location to meet floor load requirements.

About this task

Use the following steps to ensure that your location meets the floor load requirements and to determine the weight distribution area that is required for the floor load.

Procedure

1. Find out the floor load rating of the location where you plan to install the storage system.

Important: If you do not know or are not certain about the floor load rating of the installation site, be sure to check with the building engineer or another appropriate person.

2. Determine whether the floor load rating of the location meets the following requirements:
 - The minimum floor load rating that is used by IBM is 342 kg per m² (70 lb. per ft²).
 - When you install a storage system, which includes both base models and expansion models, the minimum floor load rating is 361 kg per m² (74 lb. per ft²). At 342 kg per m² (70 lb per ft²), the side dimension for the weight distribution area exceeds the 76.2 cm (30 in.) allowed maximum.
 - The per caster transferred weight to a raised floor tile is 450 kg (1000 lb.).
3. Using the following table, complete these steps for each storage system.
 - a. Find the rows that are associated with the storage system.
 - b. Locate the configuration row that corresponds with the floor load rating of the site.
 - c. Identify the weight distribution area that is needed for that storage system and floor load rating.

Note: Consult a structural engineer if you are unsure about the correct placement and weight distribution areas for your storage system.

Table 68. Floor load ratings and required weight-distribution areas

Configuration ₁	Total weight of configuration ₂	Floor load rating, kg per m ² (lb per ft ²)	Weight distribution areas _{3, 4, 5}		
			Sides cm (in.)	Front cm (in.)	Rear cm (in.)
Model 961 (2-core)	1172 kg (2585 lb)	610 (125)	2.54 (1)	76.2 (30)	76.2 (30)
		488 (100)	17.8 (7)	76.2 (30)	76.2 (30)
		439 (90)	25.4 (10)	76.2 (30)	76.2 (30)
		342 (70)	55.9 (22)	76.2 (30)	76.2 (30)
Model 961 (4-core, 8-core, 16-core, with four High-Performance Flash Enclosures)	1315kg (2900 lb)	610 (125)	5.08 (2)	76.2 (30)	76.2 (30)
		488 (100)	20.3 (8)	76.2 (30)	76.2 (30)
		439 (90)	30.5 (12)	76.2 (30)	76.2 (30)
		342 (70)	68.6 (24)	76.2 (30)	76.2 (30)
Model 961 and one 96E expansion model	2555 kg (5633 lb)	610 (125)	5.08 (2)	76.2 (30)	76.2 (30)
		488 (100)	33.0 (13)	76.2 (30)	76.2 (30)
		439 (90)	50.8 (20)	76.2 (30)	76.2 (30)
		342 (70)	76.2 (30)	76.2 (30)	76.2 (30)

Table 68. Floor load ratings and required weight-distribution areas (continued)

Configuration ₁	Total weight of configuration ₂	Floor load rating, kg per m ² (lb per ft ²)	Weight distribution areas _{3, 4, 5}		
			Sides cm (in.)	Front cm (in.)	Rear cm (in.)
Model 961 and two 96E expansion models	3776 kg (8315 lb)	610 (125)	15.2 (0)	76.2 (30)	76.2 (30)
		488 (100)	40.6 (16)	76.2 (30)	76.2 (30)
		464 (95)	53.3 (21)	76.2 (30)	76.2 (30)
Model 961 and three 96E expansion models	5080 kg (11185 lb)	610 (125)	2.54 (1)	76.2 (30)	76.2 (30)
		488 (100)	50.8 (20)	76.2 (30)	76.2 (30)
		464 (95)	66.0 (26)	76.2 (30)	76.2 (30)

Notes:

1. A storage system contains a base frame (model 961) and any expansion frames (model 96E) that are associated with it.
2. The base frame attaches to expansion frame. The expansion frame weighs 1160 kg (2550 lb) fully populated and 2375 (5225 lb) combined with the base frame. The storage enclosures weigh 22.324 kg (49.216 lb) fully configured (24 disk drives), and 14.308 kg (31.544 lb) with cables and rails and no disk drives.
3. Weight distribution areas cannot overlap.
4. Weight distribution areas are calculated for maximum weight of the frames.
Note: Keep any future upgrades in mind, and plan for the highest possible weight distribution.
5. The base and expansion frames in each storage system are bolted to each other with 5-cm (2-in.) spacers. Move one side cover and mounting brackets from the base frame to the side of the expansion frame. Side clearance for frames that are bolted together applies to both sides of the assembled frames.

Calculating space requirements

When you are planning the installation site, you must first calculate the total amount of space that is required for the storage system. Consider future expansion, and plan accordingly.

Procedure

Complete the following steps to calculate the amount of space that is required for your storage system.

1. Determine the dimensions of each frame configuration in your storage system.
2. Calculate the total area that is needed for frame configuration by adding the weight distribution area to the dimensions determined by using the table in "Meeting floor load requirements" on page 149.
3. Determine the total space that is needed for the storage system by planning the placement of each frame configuration in the storage system and how much area each configuration requires based on step 2.
4. Verify that the planned space and layout meet the service clearance requirements for each frame and storage system.

Dimensions and weight of individual models

When you are planning the floor and space requirements for your storage system, consider the dimensions and weights of the frames that compose your storage system.

The following table provides the dimensions and weights of model 961 and 96E.

Table 69. DS8870 dimensions and weights

Model	Dimensions ¹	Maximum weight of fully configured base and first expansion frames ^{2, 3}	Maximum weight of second and third expansion frames ^{2, 4}
DS8870 All-Flash Model 961	Height 193.4 cm (76 in.) Width 84.8 cm (33.4 in.) Depth 122.7 cm (48.3 in.)	1258 kg (2775 lb)	N/A
DS8870 Model 961	Height 193.4 cm (76 in.) Width 84.8 cm (33.4 in.) Depth 122.7 cm (48.3 in.)	1315 kg (2900 lb)	N/A
DS8870 Model 96E	Height 193.4 cm (76 in.) Width 84.8 cm (33.4 in.) Depth 122.7 cm (48.3 in.)	1259 kg (2820 lb)	1068 kg (2355 lb)
Notes: 1. These dimensions include casters and covers. The casters are recessed and do not require extra clearance. 2. Weight is in kilograms (kg) and pounds (lb). 3. Use this column for all base frames and for an expansion frame that can be fully configured with I/O enclosures and adapters. Expansion frames can be fully configured only when they are attached to a base frame (model 961). 4. Use this column for the second and third expansion frames that are attached to a base frame (model 961).			

Service clearance requirements

The service clearance area is the area around the storage system that IBM service representatives need to service the system.

CAUTION:

Servicing of this product or unit is to be performed by trained personnel only. (C032)

For DS8000 series, IBM services representatives must open the front and rear covers to service the storage system. Use the following minimum service clearances, which are illustrated in Figure 18 on page 153.

- For the front of the storage system, allow a minimum of 121.9 cm (48 in.) for the service clearance.
- For the rear of the storage system, allow a minimum of 76.2 cm (30 in.) for the service clearance.
- For the sides of the storage system, allow a minimum of 12.7 cm (5 in.) for the service clearance.

Unlike weight distribution areas that are required to handle floor loading, keep in mind that service clearances of adjacent unrelated storage systems can overlap.

Note: The terms *service clearance* and *weight distribution area* are often confused with each other. The service clearance is the area that is required to open the service covers and to pull out components for servicing. The weight distribution area is

the area that is required to distribute the weight of the storage system.




Figure 18. Service clearance requirements

Earthquake resistance kit installation preparation

Before an IBM service representative can install the earthquake resistance kit on any frames in your storage system, you must purchase fastening hardware and prepare the location where the kit is to be installed.

The required tasks that you must perform before the earthquake resistance kit installation depends on whether your storage system sits on a raised or a nonraised floor. For either type of installation, work with a consultant or structural engineer to ensure that your site preparations meet the requirements.

The following list provides an overview of the preparations necessary for each type of floor:

Raised floor

- Cut the necessary holes and cable cutouts in the raised floor.
- Purchase and install eyebolt fasteners in the concrete floor.

Nonraised floor

Purchase and install fasteners in the concrete floor.

Further instructions for the preparation of your site for the earthquake resistance kit are provided in “Preparing a raised floor for the earthquake resistance kit installation” and “Preparing a nonraised floor for the earthquake resistance kit” on page 159.

Preparing a raised floor for the earthquake resistance kit installation

You must prepare a raised floor and the concrete floor underneath before an earthquake resistance kit can be installed on any frame in your storage system.

Before you begin

To ensure that you meet all site requirements, obtain the service of a qualified consultant or structural engineer to help you prepare the floor.

Note: The DS8000 series supports two versions of the earthquake resistance kit. The latest version (or version 2) includes a pink eyebolt and is required for acceptance criteria for AC156 seismic certification.

Earthquake resistance kit - version 2: About this task

Figure 19 illustrates the earthquake resistance kit **1** after it is installed by IBM service representative on a raised floor.




Figure 19. Earthquake resistance kit installed on a raised floor

Complete the following steps to prepare your raised floor (see **2** in Figure 19).

Procedure

- Cut the following openings in the raised floor for each frame that uses an earthquake resistance kit:
 - Four holes for the rubber bushings of the kit to fit through the floor.
 - One cable cutout for power and other cables that connect to the rack.

Use Figure 20 as a guide for the location and dimensions of these openings. The pattern repeats for up to four frames. Dimensions are in millimeters (inches).

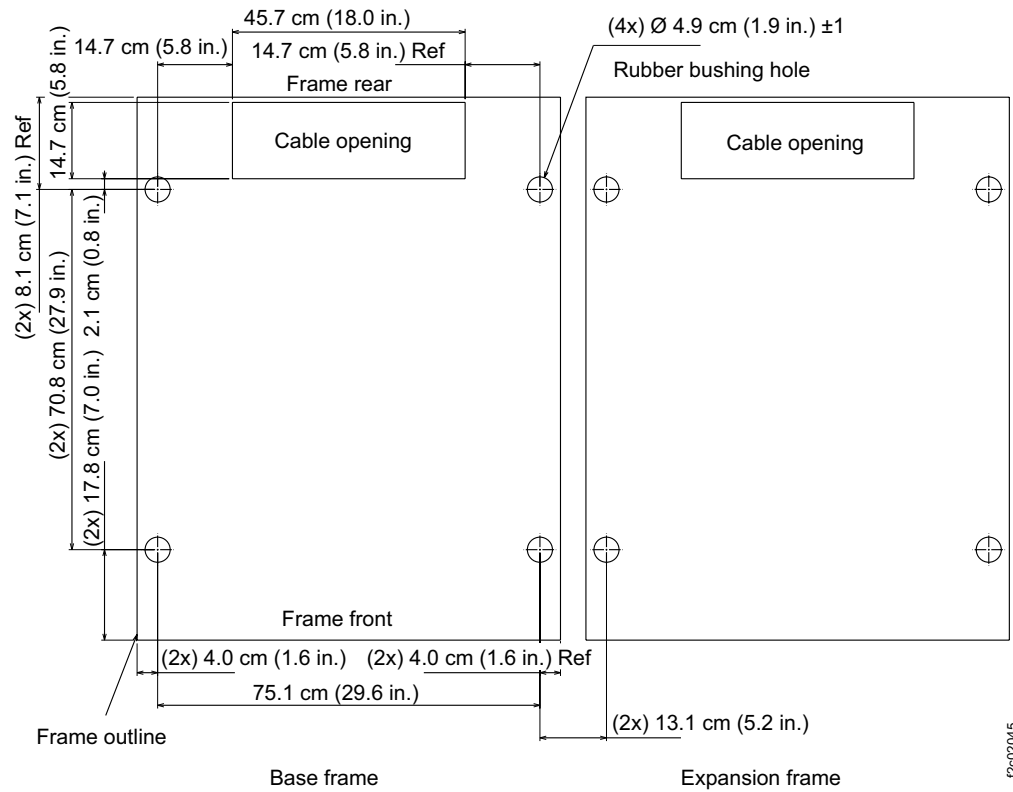


Figure 20. Locations for the cable cutouts, rubber bushing holes on raised floors, and eyebolt on concrete floors

2. Obtain eight fasteners that are heavy-duty concrete or slab floor eyebolts. These eyebolts are used to secure the earthquake resistance kit. Work with your consultant or structural engineer to determine the correct eyebolts to use, but each eyebolt must meet the following specifications.
 - Each eyebolt must withstand a 3600-pound pull force.
 - The dimensions of the eyebolt must allow the turnbuckle lower jaw of the kit to fit over the eyebolt and allow the spacer of earthquake resistance kit to fit inside the eye. See Figure 21 on page 156.




Figure 21. Eyebolt and spacer of the Earthquake Resistance Kit.

3. Install the eyebolt fasteners in the concrete floor by using the following guidelines:
 - See Figure 20 on page 155 to determine the placement of the eyebolts. The eyebolts must be installed so that they are directly below the holes that you cut in the raised floor for the rubber bushings.
 - Ensure that the installed eyebolts do not exceed a height of 10.1 cm (4 in.) from the floor to the center of the eye. This maximum height helps to reduce any bending of the eyebolt shaft.
 - Ensure that the installation allows the eyebolts to meet the required pull force after they are installed (3600-pound pull force for raised floor eyebolts).
 - If you use a threaded eyebolt that secures into a threaded insert in the floor, consider using a jam nut and washer on the shaft of the eyebolt. Talk to your consultant or structural engineer to determine whether a jam nut is necessary.

Earthquake resistance kit - version 1:

About this task

Figure 22 on page 157 illustrates the hardware kit **1** after it is installed by IBM service representative on a raised floor.




Figure 22. Earthquake resistance kit, installed on a raised floor

Use the following steps to prepare your raised floor (see **2** in Figure 19 on page 154).

Procedure

1. Cut the following openings in the raised floor for each frame that uses an earthquake resistance kit:
 - Four holes for the rubber bushings of the kit to fit through the floor
 - One cable cutout for power and other cables that connect to the rack

See Figure 23 on page 158 as a guide for the location and dimensions of these openings. The pattern repeats for up to four frames. Dimensions are in millimeters (inches).

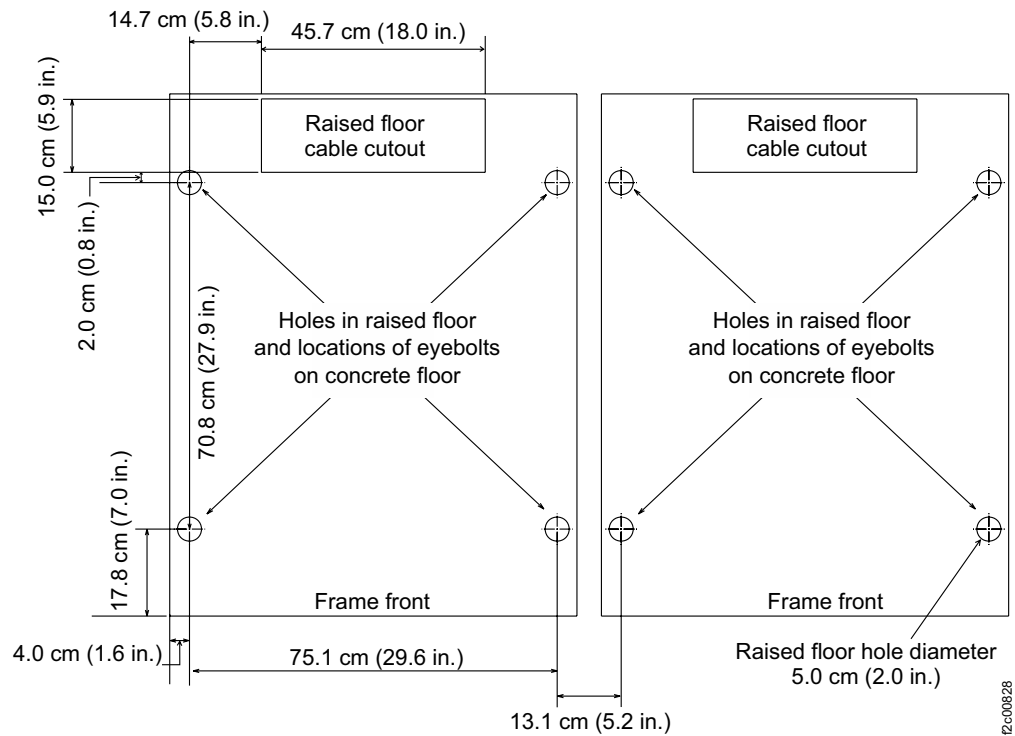


Figure 23. Locations for the cable cutouts and rubber bushing holes in the raised floor and the eyebolt installation on the concrete floor.

2. Obtain eight fasteners that are heavy-duty concrete or slab floor eyebolts. These eyebolts are used to secure the earthquake resistance kit. Work with your consultant or structural engineer to determine the correct eyebolts to use, but each eyebolt must meet the following specifications.
 - Each eyebolt must withstand a 3600-pound pull force.
 - The dimensions of the eyebolt must allow the turnbuckle lower jaw of the kit to fit over the eyebolt (**1** on Figure 24 on page 159) and allow the spacer of earthquake resistance kit to fit inside the eye (**2** on Figure 24 on page 159).




Figure 24. Eyebolt required dimensions. Dimensions are in millimeters (inches).

3. Install the eyebolt fasteners in the concrete floor by using the following guidelines:
 - See Figure 23 on page 158 to determine the placement of the eyebolts. The eyebolts must be installed so that they are directly below the holes that you cut in the raised floor for the rubber bushings.
 - Ensure that the installed eyebolts do not exceed a height of 10.1 cm (4 in.) from the floor to the center of the eye. This maximum height helps to reduce any bending of the eyebolt shaft.
 - Ensure that the installation allows the eyebolts to meet the required pull force after they are installed (3600-pound pull force for raised floor eyebolts).
 - If you use a threaded eyebolt that secures into a threaded insert in the floor, consider using a jam nut and washer on the shaft of the eyebolt. Talk to your consultant or structural engineer to determine whether a jam nut is necessary.

Preparing a nonraised floor for the earthquake resistance kit

You must prepare a nonraised floor before an earthquake resistance kit can be installed on any frame in your storage system.

Before you begin

To ensure that you meet all site requirements, obtain the service of a qualified consultant or structural engineer to help you prepare the floor.

Note: The DS8000 series support two versions of the earthquake resistance kit. The latest version (version 2) includes a pink eyebolt and is required for acceptance criteria for the AC156 seismic certification.

Earthquake resistance kit version 2:

About this task

Figure 25 provides an illustration of the earthquake resistance kit (**1**) after the IBM service representative installs it on the nonraised floor.

Before the IBM service representative installs the kit, you must prepare the area that is shown as **2** in Figure 25. This figure shows two of the most common fasteners that you can use.




Figure 25. Earthquake resistance kit installed on a nonraised floor.

f2c02047

Use the following steps to prepare your nonraised floor:

Procedure

1. Obtain eight fastener sets for each frame that uses the earthquake resistance kit. These fastener sets are used to secure the earthquake resistance kit load plate. The type of fastener set that you use can be determined by your consultant or structural engineer. However, each bolt or stud must meet the following specifications:

- Each fastener set must withstand a 2400-lb pull force.
- The fasteners must have a dimension that fits into the load plate holes, which are each 2.7 cm (1.0 in.) in diameter.
- The fasteners must be long enough to extend through and securely fasten a load plate that is 3.0 cm (1.2 in.) thick. The fasteners must also be short enough so that the height of the installed fastener does not exceed 6.5 cm (2.5 in.). This maximum height ensures that the fastener can fit under the frame.

The following examples provide descriptions of nonraised floor fastener sets:

- Threaded hole insert that is secured into the concrete floor and a bolt (with a washer) that screws into the insert
- Threaded stud that is secured into the concrete floor with a nut (with a washer) that screws over it

Figure 25 on page 160 illustrates the fastener sets.

2. Work with your consultant or structural engineer and use the following guidelines to install the fasteners in the concrete floor:
 - Use Figure 26 on page 162 to determine the placement of the fasteners. The pattern repeats for up to three frames. Dimensions are in millimeters (inches).
 - Ensure that the installed fasteners do not exceed a height of 6.5 cm (2.5 in.) from the floor. This maximum height ensures that the fastener can fit under the frame.
 - Ensure that the installation allows the fasteners to meet the required pull force after they are installed (2400 lb pull force).
 - If you use a threaded bolt that secures into a threaded insert in the floor and the bolt extends longer than 3.0 cm (1.2 in.), which is the thickness of the load plate, consider using a jam nut and a washer on the shaft of the bolt so that the load plate can be secured snugly to the floor. Talk to your consultant or structural engineer to determine whether a jam nut is necessary.




Figure 26. Locations for fastener installation (nonraised floor)


- When the IBM service representative arrives to install the earthquake resistance kit, provide the other fastener parts (2) in Figure 25 on page 160 so that the representative can use these parts secure the load plates onto the floor.

Earthquake resistance kit version 1:

About this task

Figure 27 on page 163 provides an illustration of the earthquake resistance kit (1) after the IBM service representative installs it on the nonraised floor. This figure illustrates two of the most common fasteners that you can use.

Before the IBM service representative installs the kit, you must prepare the area that is shown as 3 in Figure 27 on page 163.



f2c00827

Figure 27. Earthquake resistance kit installed on a nonraised floor.

Use the following steps to prepare your nonraised floor:

Procedure

1. Obtain eight fastener sets for each frame that uses the earthquake resistance kit. These fastener sets are used to secure the earthquake resistance kit load plate. The type of fastener set that you use can be determined by your consultant or structural engineer. However, each bolt or stud must meet the following specifications:
 - Each fastener set must withstand a 2400-lb pull force.
 - The fasteners must have a dimension that fits into the load plate holes, which are each 2.7 cm (1.0 in.) in diameter.
 - The fasteners must be long enough to extend through and securely fasten a load plate that is 3.0 cm (1.2 in.) thick. The fasteners must also be short enough so that the height of the installed fastener does not exceed 6.5 cm (2.5 in.). This maximum height ensures that the fastener can fit under the rack.

The following examples provide descriptions of nonraised floor fastener sets:

- Threaded hole insert that is secured into the concrete floor and a bolt (with a washer) that screws into the insert

- Threaded stud that is secured into the concrete floor with a nut (with a washer) that screws over it

Figure 27 on page 163 illustrates the fastener sets. The pattern repeats for up to three frames. Dimensions are in millimeters (inches).

2. Work with your consultant or structural engineer and use the following guidelines to install the fasteners in the concrete floor:
 - Use Figure 28 to determine the placement of the fasteners.
 - Ensure that the installed fasteners do not exceed a height of 6.5 cm (2.5 in.) from the floor. This maximum height ensures that the fastener can fit under the rack.
 - Ensure that the installation allows the fasteners to meet the required pull force after they are installed (2400 lb pull force).
 - If you use a threaded bolt that secures into a threaded insert in the floor and the bolt extends longer than 3.0 cm (1.2 in.), which is the thickness of the load plate, consider using a jam nut and a washer on the shaft of the bolt so that the load plate can be secured snugly to the floor. Talk to your consultant or structural engineer to determine whether a jam nut is necessary.




Figure 28. Locations for fastener installation (nonraised floor).

3. When the IBM service representative arrives to install the earthquake resistance kit, provide the other fastener parts (2 in Figure 27 on page 163) so that the representative can use these parts secure the load plates onto the floor.

Planning for power requirements

You must select a storage system location that meets specific power requirements.

When you consider the storage system location, consider the following issues:

- Power control selections
- Power outlet requirements
- Input voltage requirements
- Power connector requirements
- Remote force power off switch requirements
- Power consumption and environment

IBM cannot install the storage system if your site does not meet these power requirements.

Attention: Implementation of surge protection for electronic devices as described in the EN 62305 standard or IEEE Emerald Book is recommended. If a lightning surge or other facility transient voltages occurs, a surge-protection device limits the surge voltage that is applied at the storage system power input. A surge-protection device is required for facilities in Korea or customers that conform to the European EMC Directive or CISPR 24.

Overview of storage system power controls

The storage system contains power controls on the frames. The power controls can be configured by an IBM service representative. The power controls can also be accessed through the management console.

The storage system has the following manual power controls in the form of physical switches that are on the racks:

Local/remote switch

(Available on base frames) The local/remote switch setting determines your use of local or remote power controls. When you set the switch to local, the local power on/local force power off switch controls power in the storage system. You can access this switch by opening the rear cover of the storage system. When the local/remote switch is set to remote, the power for the storage system is controlled by remote power control settings that are entered in the DS8000 Storage Management GUI or DS Service GUI.

Planning requirements: None.

Local power on/local force power off switch

(Available on base frames) The local power on/local force power off switch initiates a storage system power-on sequence or a storage system force power off sequence. This switch is applicable only when the local/remote switch is set to local. You can access this switch by opening the rear cover of the storage system.

Planning requirements: None.

Emergency power off switch

(Available on all frames) If activated, the emergency power-off switch causes the individual frame to immediately drop all power, including any power that is provided by the battery system. When activated, this switch overrides all other power controls for the specific rack. This switch is located behind the covers of each frame.

Attention: Use this switch only in extreme emergencies. Using this switch might result in data loss.

Planning requirements: None.

The following power controls can be configured by an IBM service representative. You can also use the following power controls through the DS8000 Storage Management GUI (running on the management console):

Local power control mode

(Visible in the DS8000 Storage Management GUI) You cannot change this setting in the DS8000 Storage Management GUI. This mode is enabled when the local/remote switch that is on the storage system is in the local position. When this setting is used, the local power on/local force power-off switch that is on the storage system controls the power.

Planning requirements: None.

Remote power control mode

(Visible in the DS8000 Storage Management GUI) If you select the **Remote** power control mode, you choose one of the following remote mode options.

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

Remote Management Console, Manual

Your use of the DS8000 Storage Management GUI power on/off page controls when the storage system powers on and off.

Remote Management Console, Scheduled

A schedule, which you set up, controls when the storage system powers on and off.

Remote Management Console, Auto

This setting applies only in situations in which input power is lost. In those situations, the storage system powers on as soon as external power becomes available again.

Remote Auto/Scheduled

A schedule, which you set up, controls when the storage system powers on and off. A power-on sequence is also initiated if the storage system was powered off due to an external power loss while the storage systems are scheduled to be on and external power becomes available again.

Remote zSeries Power Control

One or more attached z Systems or S/390 hosts can control the power-on and power-off sequences.

Power outlet requirements

Plan for the required power outlets for the installation of your storage system.

The following power outlets are required:

- Two independent power outlets for the two power cords that are needed by each base and expansion frames.

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

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

Input voltage requirements

When you plan for the power requirements of the storage system, consider the input voltage requirements.

The following figure provides the input voltages and frequencies that the storage system supports.

Table 70. Input voltages and frequencies

Characteristic	Low voltage (three-phase or single-phase)	High voltage (three-phase)
Nominal input voltages	200, 208, 220, or 240 RMS V ac	380, 400, or 415 RMS V ac
Minimum input voltage	180 RMS V ac	315 RMS V ac
Maximum input voltage	256 RMS V ac	456 RMS V ac
Customer wall breaker rating (1 ph, 3 ph)	50-60 Amps	30-35 Amps
Steady-state input frequencies	50 ± 3 or 60 ± 3.0 Hz	50 ± 3 or 60 ± 3.0 Hz
PLD input frequencies (<10 seconds)	50 ± 3 or 60 ± 3.0 Hz	50 ± 3 or 60 ± 3.0 Hz

Power connector requirements

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

Table 71 on page 168 shows the power cords and the inline connectors and receptacles types that they support. Find the power cord row that is appropriate for your site. Ensure that your site meets the specified power connector requirements.

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

Attention:

- For reliable operation do not use Ground Fault Circuit Interrupter (GFCI), Earth Leakage Circuit Breaker (ELCB), or Residual Current Circuit Breaker (RCCB) type circuit breakers with the DS8870. The DS8870 is certified for safe operation and compliant with IEC, EN, UL, and CSA 60950-1 standards. If local electrical practice requires leakage detection circuit breakers, they must be rated at minimum to 300 mA or larger to reduce the risk of outage due to spurious actuation.
- Low-voltage, three-phase installations (200-240 V) require wall circuit breakers that have a rating of 50 to 60 A.
- High-voltage, three-phase installations (380-415 V) require wall circuit breakers that have a rating of 30 to 35 A.

Note: Special applications that use 30 A power cords must rely on 30 A wall circuit breakers for power cord protection.

- Low voltage, single-phase, and three-phase low voltage installations require wall circuit breakers that have a rating of 50 to 60 A.
- High voltage, three-phase high-voltage installations require wall circuit breakers that have a rating of 30 to 35 A.
- Do not exceed the wire rating of the facility and ensure that separate protected branch circuits are used for each cord in planning for redundancy.

Table 71. DS8870 power cords

Power cord feature code	Power cord description	Inline connector	Receptacle
1061 _{1, 2, 3, 4, 5}	Single-phase power cord, 200V - 240V, 60A, 3-pin connector	HBL360C6W, Pin and Sleeve Connector, IEC 309, 2P3W	HBL360R6W, ac Receptacle, IEC 60309, 2P3W
1068 _{2, 3, 4}	Single-phase power cord, 200V - 240V, 63A, no connector	Not applicable	Not applicable
1072 _{1, 2, 3, 4, 5}	Top exit single-phase power cord, 200V - 240V, 60A, 3-pin connector	HBL360C6W, Pin and Sleeve Connector, IEC 309, 2P3W	HBL360R6W, ac Receptacle, IEC 60309, 2P3W
1073 _{2, 3, 4}	Top exit single-phase power cord, 200V - 240V, 63A, no connector	Not applicable	Not applicable
1080 _{1, 2, 3, 4, 5, 6}	Three-phase power cord, high voltage (five-wire 3P+N+G) 380-415V (nominal), 30A, IEC 60309, 5-pin customer connector	HBL530C6V02, Pin and Sleeve Connector, IEC 60309, 4P5W	HBL530R6V02, AC Receptacle, IEC 60309, 4P5W
1081 _{2, 3, 4}	Three-phase high voltage (five-wire 3P+N+G), 380-415V, 32A, no customer connector provided	Not applicable	Not applicable
1082 _{1, 2, 3, 4, 5}	Three-phase power cord, low voltage, 200V - 240V, 60A, 4-pin connector	HBL460C9W, Pin and Sleeve Connector, IEC 309, 3P4W	HBL460R9W, ac Receptacle, IEC 60309, 3P4W
1083 _{2, 3, 4}	Three-phase high voltage (five-wire 3P+N+G), 380-415V, 32A, no customer connector provided	Not applicable	Not applicable

Table 71. DS8870 power cords (continued)

Power cord feature code	Power cord description	Inline connector	Receptacle
1084 _{1, 2, 3, 4, 5}	Three-phase low voltage (four-wire 3P+G), 200-240V, 60A, IEC 60309-4-pin customer connector	HBL460C9W, Pin and Sleeve Connector, IEC 309, 3P4W	HBL460R9W, ac Receptacle, IEC 60309, 3P4W
1085 _{1, 2, 3, 4, 5, 6}	Top exit Three-phase high voltage (five-wire 3P+N+G), 380-415V (nominal), 30A, IEC 60309, 5-pin customer connector	HBL530C6V02, Pin and Sleeve Connector, IEC 60309, 4P5W	HBL530R6V02, AC Receptacle, IEC 60309, 4P5W

Notes:

1. The customer connector must be IEC 60309-compliant. Power cords that are used with DS8100, DS8300, DS8700, and DS8800 cannot be used again with the DS8870.
2. Power cords without inline connectors are rated at 600 V ac. Other power cords are rated as described in the table. Single-phase power cord has 2W+G configuration. Low voltage three-phase power cord has 3W+G configuration. High-voltage three-phase power cord has a 3W+N+G configuration.
3. The conductor size for three-phase power cords is 6 mm² (10awg). The conductor size for single-phase power cords is 10 mm² (6awg).
4. Power cords that exit the bottom are 4.2 M (14 ft) from the lowest point where they exit the frame to the mating face of the plug or bare leads. Power cords that exit the top are 1.8 M (6 ft) from the highest point from the frame to the mating face of the plug or bare leads.
5. The IEC60309 receptacle must be installed in a metal-backed box with the green wire ground-connected to the grounding lug within the box. Ensure continuity between the box and the metallic shielding of the liquid-tight conduit.
6. Feature codes 1080 and 1085 are supported in North America.

The following three figures show the plug (connector on the line cord).




Figure 29. Single-phase, 200 - 240 V, 60A, IEC60309 cable connector pins




Figure 30. Three-phase low voltage (200 - 240 V) 60A, IEC60309 cable connector pins




Figure 31. Three-phase high voltage (380 - 415 V) 30A, IEC60309 cable connector pins

The following three figures show the receptacle or inline connector (the power source, a mirror image of the plug).




Figure 32. Single-phase, 200 - 240 V, 60A, IEC60309 customer outlet pins




Figure 33. Three-phase low voltage (200 - 240 V) phase-to-phase, 60A, IEC60309 customer outlet pins




Figure 34. Three-phase high voltage (380 - 415 V) phase-to-phase, 30A, IEC60309 customer outlet pins

Power consumption and environmental information

When you are planning to meet the power requirements for the DS8000 series, consider the power consumption and other environmental points of the storage system.

This power consumption and environmental information for the IBM DS8870 is provided in Table 72.

Table 72. Power consumption and environmental information for models 961 and 96E

Measurement	Unit of measure	Base frame (model 961)	Expansion frame (model 96E)
Peak electric power ^{1, 3}	kilovolt amperes (kVA)	7.3	5.8
Thermal load	British thermal units (BTU) per hour	25,000	19,605
Capacity of exhaust	cubic meters per minute (cubic feet per minute or CFM)	44.2 (1500)	51.8 (1800)
Ground leakage current	milliamperes (mA)	60	60
Startup current	amperes (A or amp)	≤ 100	≤ 100
Startup current duration	microseconds (μs or μsec)	< 200	< 200
Idle and operating sound power level, LWAd ²	A-weighted bels (B)	7.5	7.5

Notes:

- The values represent data that was obtained from typical systems, which are configured as follows:
 - Standard base frames that contain 15 disk drive sets (16 drives per disk drive set, 15 disk drive sets x 16 = 240 disk drives) and Fibre Channel adapters.
 - All-flash configurations that contain 8 sets of fully configured high-performance storage enclosures and 16 Fibre Channel adapters.
 - Expansion models that contain 21 drive sets per storage enclosure (21 drive sets x 16 = 336 drives) and Fibre Channel adapters.
- LWAd is the statistical upper-limit A-weighted sound power level, expressed in bels, declared in conformance with ISO 9296. The values correspond to DS8870 containing 240 drives. The values represent data that is obtained during preliminary testing. Bel relates to decibels (dB) as follows: $1 B = 10 dB$. The ratings are rounded to the nearest 0.1 B. Measurements are in conformance with ISO 7779.
- All frames and configurations that are used in single-phase mode must not exceed 8 kVA.

DS8000 support appliances

The DS8000 support appliances are mounted external to the DS8000 in customer-supplied rack space. These appliances provide platforms for software such as IBM Security Key Lifecycle Manager (SKLM) isolated key-manager and Copy Services Manager (CSM). Model 961 feature code 1760, and model AP1 feature code 1761 include a 1U monitor keyboard assembly and required rack-mount line cords. Model AP1 is available in single server (feature code 1761) and dual server (feature code 1761 and 1762 ordered together) configurations.

The model AP1 tray includes a keyboard video monitor (KVM) switch to allow the dual server configuration to utilize a shared monitor. The bezel in the tray has a label with the machine type, model, and serial number. This information is needed for hardware service.

Full disk encryption offerings must be activated before use by the IBM Systems Lab Services team. See “Planning for full-disk encryption activation” on page 197 for more information.

Table 73. Appliance rack space and power characteristics

Feature code	Description	Power consumption	Input voltage
1760	1U appliance processor	< 580 watts	100 - 240 V ac
1761	2U appliance tray with one Intel 4-core processor assembly	< 100 watts	100 - 240 V ac
1762	2U appliance tray with two Intel 4-core processor assemblies	< 190 watts	100 - 240 V ac
Monitor included with feature code 1760 or 1761	1U keyboard monitor	< 240 watts	100 - 240 V ac

Acoustic declaration for the DS8000 series

Table 74 describes the acoustic declaration information for the DS8000 series.

Table 74. Acoustic declaration for fully configured DS8000 series

Model	Declared A-weighted sound power level, LWAd (B) ^{1, 4}		Declared A-weighted sound pressure level, LpAm (dB) ^{2, 3, 4}	
	Operating	Idling	Operating	Idling
Model 961	7.9	7.9	61	61
Model 96E	7.9	7.9	61	61

Notes:

- LWAd is the statistical upper-limit A-weighted sound power level (rounded to the nearest 0.1 B).
- LpAm is the mean A-weighted emission sound pressure level that is measured at the 1-meter bystander positions (rounded to the nearest dB).
- 10 dB (decibel) = 1 B (bel).
- All measurements made in conformance with ISO 7779 and declared in conformance with ISO 9296.

Planning for environmental requirements

You must install your storage system in a location that meets the operating environment requirements for correct maintenance.

Take the following steps to ensure that you meet these requirements:

1. Note where air intake locations are on the models that compose your storage system.
2. Verify that you can meet the environmental operating requirements at the air intake locations.
3. Consider optimizing the air circulation and cooling for the storage system by using a raised floor, adjusting the floor layout, and adding perforated tiles around the air intake areas.

Fans and air intake areas

The storage system provides air circulation through various fans throughout the frame. You must maintain the correct operating environment requirements for your models at each air intake location.

Table 75 summarizes fan, intake, and exhaust locations.

Table 75. Machine fan location

DS8870 Fan Location	Machine Location	Intake Location	Exhaust Location
Entire machine	Entire	Front covers	Rear covers
Power complex	Left-side	Front covers	Rear covers

Operating environment requirements

You must meet specific operating environment requirements at all the air intake locations of your storage system.

The operating points vary depending on the state of the system. The system can be in the following states:

- Powered on
- Powered off
- In storage

Powered on:

Plan for the operating ranges and recommended operating points of the storage system.

Table 76 provides the operating ranges for your storage system when its power is on.

Table 76. Operating extremes with the power on

Measurement	Value
Altitude	0 - 2133 m (0 - 7000 ft)
Dry bulb temperature	16 - 32°C (60 - 90°F)
Relative humidity	20 - 80%
Wet bulb temperature (maximum)	23°C (73°F)

Table 77 provides the optimum operating points for your storage system with its power on.

Table 77. Optimum operating points with the power on

Measurement	Value
Temperature	22°C (72°F)
Relative humidity	45%

Table 78 provides the operating ranges for a storage system with power on.

Table 78. Optimum operating ranges with power on

Measurement	Value
Temperature	16 - 32°C (60 - 90°F)
Relative humidity	40 - 50%

Powered off:

Plan for the required temperature and humidity ranges when the storage system is off.

Table 79 provides the temperatures and humidity requirements for your storage system when the power is off.

Table 79. Temperatures and humidity with the power off

Measurement	Value
Temperature	10 - 43°C (50 - 110°F)
Relative humidity	8 - 80%
Wet bulb temperature (maximum)	27°C (80°F)

In storage:

Plan for the required temperature and humidity ranges when the storage system is in storage.

Table 80 provides the temperatures and humidity requirements for storing your storage system.

Table 80. Temperatures and humidity while in storage

Measurement	Value
Temperature	1 - 60°C (34 - 140°F)
Relative humidity	5 - 80%
Wet bulb temperature (maximum)	29°C (84°F)

Corrosive gasses and particulates

Plan for air quality that meets standards for corrosive gases and particulates.

The DS8000 series is designed to operate reliably in a general business-class environment. A general business-class environment is one that has automated 24x7x365 temperature and humidity controls and also operates with G1 specifications for corrosive gases and P1 specifications for particulates.

Operating vibration requirements

The vibration levels that are designed for DS8870 comply with class V1 requirements included in the product classes for vibration.

DS8870 is designed to operate under the vibration V1 levels that are described in Table 81. Additional information includes random vibration PSD profile breakpoints and operational shock levels. See Table 82 and Table 83.

Table 81. Vibration levels for DS8870

Class	grms	g Peak Sine ^{1, 2}
V1L	0.10	0.06 @ 50 and 60 Hz
Notes:		
1. All values in this table are in g^2/Hz		
2. g is the peak g level of an approximate half-sine pulse.		

Table 82. Random vibration PSD profile breakpoints for DS8870

Class	5 Hz	17 Hz	45 Hz	48 Hz	62 Hz	65 Hz	150 Hz	200 Hz	500 Hz
V1L	2.0×10^{-7}	2.2×10^{-5}	2.2×10^{-5}	2.2×10^{-5}	2.2×10^{-5}	2.2×10^{-5}	2.2×10^{-5}	2.2×10^{-5}	2.2×10^{-5}
Note: All values in this table are in g^2/Hz .									

Table 83. Operational shock levels for DS8870

Class	Axis	g^1	pw^2
S1	Vertical	3.5	3.0
Notes:			
1. g is the peak g level of an approximate half-sine pulse.			
2. " pw " is the pulse width in milliseconds.			

Contamination information

You must consider the air quality and contamination levels at your installation site.

Airborne particulates (including metal flakes or particles) and reactive gases that act alone or in combination with other environmental factors, such as humidity or temperature, might pose a risk to the storage system hardware. Risks that are posed by the presence of excessive particulate levels or concentrations of harmful gases include damage that might cause the system to malfunction or cease functioning altogether. This specification presents limits for particulates and gases that are intended to avoid such damage. The limits must not be viewed or used as definitive limits because numerous other factors, such as temperature or moisture content of the air, can influence the impact of particulates or environmental corrosives and gaseous contaminant transfer.

Attention: In the absence of specific limits that are presented in this document, you must implement practices that maintain particulate or gas levels that are consistent with the protection of human health and safety. If IBM determines that the levels of particulates or gases in your environment damaged the storage system, the warranty is void. Implementation of correctional measures is a customer responsibility.

The following criteria must be met:

Gaseous contamination

Severity level G1 according to ANSI/ISA 71.04-1985¹, which states that the reactivity rate of copper coupons must be fewer than 300 Angstroms per month ($\text{\AA}/\text{month}$, $\approx 0.0039 \mu\text{g}/\text{cm}^2\text{-hour}$ weight gain)². In addition, the reactivity rate of silver coupons must be less than $300\text{\AA}/\text{month}$ ($\approx 0.0035 \mu\text{g}/\text{cm}^2\text{-hour}$ weight gain)³. The reactive monitoring of gaseous corrosivity is conducted approximately 2 inches (5 cm) in front of the rack on the air inlet side at one-quarter and three-quarter frame height off the floor, or where the air velocity is much higher.

Particulate contamination

Data centers must meet the cleanliness level of ISO 14644-1 class 8. For data centers without airside economizers, the ISO 14644-1 class 8 cleanliness can be met by selecting one of the following filtration methods:

- The room air can be continuously filtered with MERV 8 filters.
- Air entering a data center can be filtered with MERV 11, or preferably MERV 13 filters.

For data centers with airside economizers, the choice of filters to achieve ISO class 8 cleanliness depends on the specific conditions present at that data center. The deliquescent relative humidity of the particulate contamination must be more than 60% RH⁴. Data centers must be free of zinc whiskers⁵.

1. ANSI/ISA-71.04.1985. *Environmental conditions for process measurement and control systems: Airborne contaminants*. Instrument Society of America, Research Triangle Park, NC, 1985.
2. The derivation of the equivalence between the rate of copper corrosion product thickness growth in $\text{\AA}/\text{month}$ and the rate of weight gain assumes that Cu_2S and Cu_2O grow in equal proportions.
3. The derivation of the equivalence between the rate of silver corrosion product thickness growth in $\text{\AA}/\text{month}$ and the rate of weight gain assumes that Ag_2S is the only corrosion product.
4. The deliquescent relative humidity of particulate contamination is the relative humidity at which the dust absorbs enough water to become wet and promote corrosion, ion migration, or both.
5. Surface debris is randomly collected from 10 areas of the data center on a 1.5 cm diameter disk of sticky, electrically conductive tape on a metal stub. If examination of the sticky tape in a scanning electron microscope reveals no zinc whiskers, the data center is considered free of zinc whiskers.

Cooling the storage complex

You can take steps to optimize the air circulation and cooling for your storage systems.

About this task

To optimize the cooling around your storage systems, prepare the location of your storage systems as recommended in the following steps.


Note: The installation of a storage system is done by IBM service representatives. However, the following steps describe the process that is needed to optimize the air circulation and cooling for your storage systems.

DS8870


1. Prepare for the installation of the storage system on a raised floor. Although the storage system can be installed on a nonraised floor, installing the storage system on a raised floor provides increased air circulation for better cooling.
2. Install perforated tiles in the front of each base frame and expansion frame as follows:
 - a. For a stand-alone base frames, install two fully perforated tiles in front of each base frame, as shown in Figure 35 on page 178 in the single-machine examples (**1** and **2** in the figure).
 - b. For a row of storage systems, install a row of perforated tiles in front of the storage systems as shown in Figure 35 on page 178 (**3** and **4** in the figure).
 - c. For groupings of storage systems, where a hot aisle/cold aisle layout is used, use a cold aisle row of perforated tiles in front of all storage systems. A possible minimal configuration is shown in Figure 35 on page 178 (**5** in the figure).

Note: Keep in mind that the dimensions represented in Figure 35 on page 178 might not be adequate for floor-load requirements.

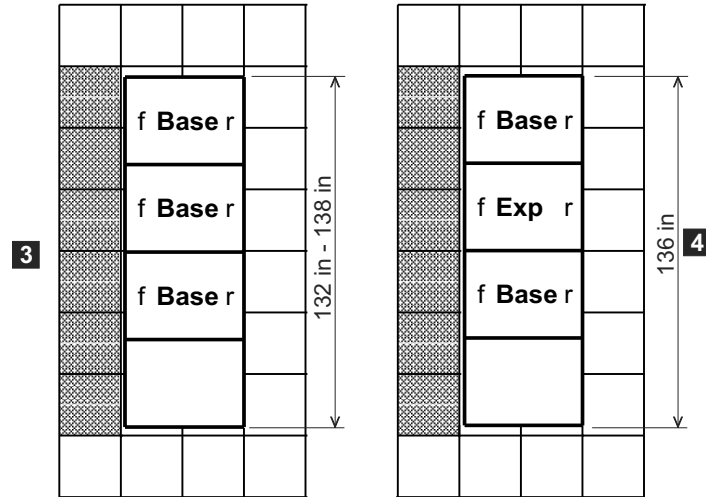
f = front Base = Base model
 r = rear Exp = Expansion model



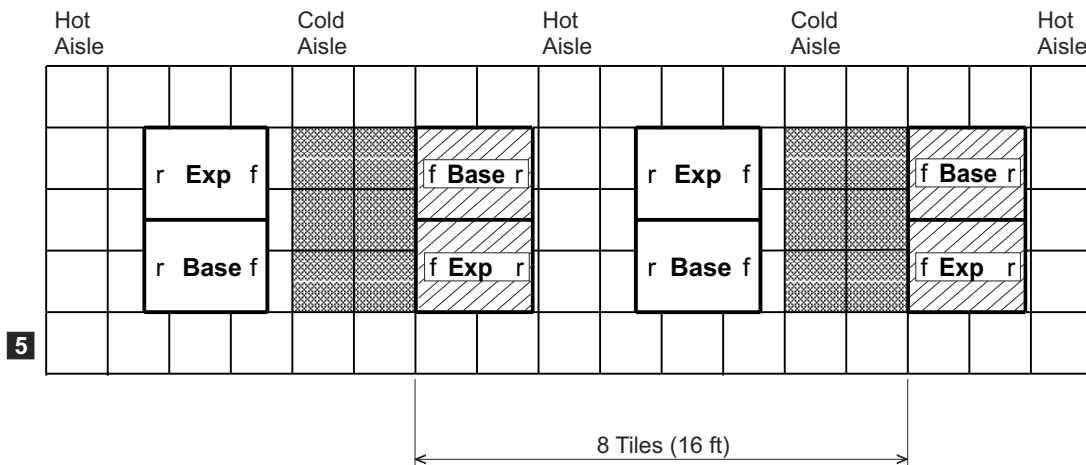
Single Machine - 2 examples



Row of Machines - 2 examples



Hot Aisle / Cold Aisle installation



f2c01589

Figure 35. DS8870 layouts and tile setup for cooling

Planning for safety

You must consider various safety issues when you plan your storage system location.

The following list identifies some of the safety issues you must consider:

- Fire suppression
- Earthquake safety

Providing a fire-suppression system

Set up an environment that supports the temperature, cooling, and operating requirements of your storage system. You are responsible for providing a fire suppression system for your storage system.

About this task

IBM designs and manufactures equipment to internal and external standards that require certain environments for reliable operation. Because IBM does not test any equipment for compatibility with fire-suppression systems, IBM does not make compatibility claims of any kind. IBM does not provide recommendations on fire-suppression systems.

Procedure

1. Consult your insurance underwriter, local fire marshal, or local building inspector about selecting a fire-suppression system that provides the correct level of coverage and protection.
2. Set up an environment that supports the temperature and cooling requirements for your storage system as described in the environmental temperature requirements-planning area.

Earthquake preparedness alternatives

If you are installing your storage system in an area that is prone to earthquakes, plan for special installation methods to minimize earthquake damage to your system.

An unsecured base or expansion frame can topple or be displaced during an earthquake. This displacement puts both the storage system and your personnel in danger. To help prevent damage, restrain your storage system by using one of the following two methods:

Restraint method

Allows some system movement and provides for both personnel safety and protection of your storage system. The DS8000 earthquake resistance kit provides this type of earthquake protection.

Hard mounting

Physically attaches your system to the floor. This method increases the safety of personnel during an earthquake. However, the method damages the storage system because the system absorbs most of the shock. IBM does not support hard mounting.

Planning for external management console installation

If you use an external management console (feature code 1130), you must plan for its installation.

You must provide the following for the installation of the external management console:

- A location for the external management console that allows the Ethernet cable, which is 31 m (101 ft), to reach the management console on one end and the storage system on the other end. You can locate your management console farther away, if necessary; however, this location can make servicing your storage system more difficult, and you must provide two of your own Ethernet cables. Any Ethernet cables that you supply must meet the following requirements:
 - The cables must be straight 4-pair UTP CAT 5E (or better) cables with RJ-45 jacks.
 - The length of each cable cannot exceed 100 m (328 ft).

- The two private network cables are labeled or marked on both ends to make it obvious which one belongs to the gray network and, which one belongs to the black network.
- A rack in which to mount the management console. This rack can be an IBM rack or a rack from another company, but it must meet specifications.
- An external management console must be installed in an environment that does not exceed 35 degrees Celsius (°C).
- Two outlets that are within 3.1 m (10 ft) of the external management console. Typically, these outlets are in the rack that you provide. The outlets are used by the country-specific power cords that you select when you order the external management console.

Both the internal and external management consoles are dedicated servers. The first management console is always located inside the storage system.

Rack specifications for the external management consoles

Use rack specifications information to ensure that you have the required environment for a non-IBM rack if you plan to install your external management console in a non-IBM rack.

If you install your external management console in a non-IBM rack, ensure that the rack satisfies the required specifications. The following specifications are required for a non-IBM rack:

- The rack must meet EIA-310-D standards for mounting flanges and hole locations.
- The front-to-rear distance of the mounting flanges must be 720 mm (28.3 in.).
- The thickness of the mounting flanges must be between 1.9 and 3.3 mm (0.08 and 1.3 in.).
- The mounting flanges must have either 7.1 mm (0.28 in.) diameter holes or 9.6 mm (0.38 in.) square holes on the standard EIA hole spacing.
- The rack must have a minimum depth of 70 mm (2.76 in.) between the front mounting flange and inside of the front door for appropriate cooling.
- The rack must have a minimum depth of 157 mm (6.2 in.) between the rear mounting flange and inside of the rear door to install the external management console and provide cable management space.
- The minimum side-to-side clearance in the rack between the front and rear mounting flanges must be of 467 mm (18.2 in.) to accommodate the width of the server and the slide-mounting brackets.
- The front edge of the keyboard/display tray is 19 in. wide including the tabs for the front thumbscrews. To use the tray, it must be extended out fully on its rails. If the rack is equipped with front covers, ensure that they are hinged in such a way that in their fully open position they allow the keyboard/display tray to be fully extended.
- The rack must include perforated front and rear doors and must not prevent the flow of cool air into or out of the rack.
- The weight-handling capacity of the rack must be able to support the maximum rack configuration, including all servers, external cables, and PDUs.
- The rack must provide correct stabilization so that the rack does not become unstable when the external management console or other equipment are pulled out for use or for service.

Planning for network and communications requirements

You must locate your storage systems in a location that meets the network and communications requirements.

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

- Management console network requirements
- Remote power control requirements
- Host attachment requirements
- SAN considerations

Management console network requirements

You must plan for the network requirements of the management console.

Each management console requires a dedicated connection to the network.

Note: If you plan on accessing the DS CLI or DS8000 Storage Management GUI and have a firewall between the management console and your network, open the following TCP/IP ports before installation: 1750, 1751, 1755 for CLI, and 8452 for the DS8000 Storage Management GUI .

Remote support connection requirements

IBM highly recommends remote-support connections so that storage system can use the call home feature to report failures and transmit pertinent debugging data to IBM. When call home is used, your IBM service representative can quickly isolate and resolve potential issues with the storage system.

The strategic and preferred remote support connectivity method is Internet SSL (Secure Socket Layer) for management-console-to-IBM communication, and Assist On-site (AOS) for IBM remote access to the management console and the storage system. AOS provides a network-type connection that is secured by SSL and state-of-the-art encryption technology. AOS is installed on a PC that is provided and maintained by the customer. Alternatively, if the customer's security guidelines allow, AOS can be turned on and configured on the management console. Contact your IBM service representative for more details.

Another legacy remote support solution is the management console VPN. The VPN is used for outbound and inbound connectivity. Because the VPN is always initiated by the management console, either the local management-console service interface or DS CLI must be used to start the VPN. If a modem is configured at the same time, it can be used by IBM service representative to initiate the VPN connection. All IBM remote support solutions provide you with an interface to control and secure access to the storage system. IBM also provides activity and authentication logging.

The modem continues to be an option for inbound and outbound remote support connectivity. In conjunction with the modem, Internet FTP (File Transfer Protocol) can be configured for faster offload of memory dump and trace data. To support the modem connection, you must provide the following equipment close enough to each management console:

- One analog telephone line per management console for initial setup
- A telephone cable to connect the modem to a telephone jack

Note that the network to support modem remote connection may be discontinued in future versions of the DS8000 product. Consider using VPN or AOS for remote

connection to avoid future disruption. All IBM remote support solutions provide an interface to control and secure access to the storage system. IBM also provides activity and authentication logging.

The remote support option you chose must be ready at the time of the initial installation of the DS8000. This includes opening your firewalls for Internet based remote support.

Remote power control requirements

Use the remote power control settings to control the power of your storage complex. Settings can be controlled through the DS8000 Storage Management GUI running on the management console.

There are several settings for remote power control. Only the remote zSeries power control setting requires planning on your part.

The remote zSeries power control setting allows one or more attached z Systems or S/390 hosts to control the power-on and power-off sequences for your storage system. If you use the remote zSeries power control setting, you must meet the following requirements:

- Order the remote zSeries power control feature, FC1000.
- Allow up to four interfaces for remote zSeries power control.

Host attachment communication requirements

Use host attachment communication requirements information to connect the host attachments in your network.

- You must use worldwide port names to uniquely identify Fibre Channel adapters that are installed in your host system.
- For open-system hosts with Fibre Channel adapters, keep in mind that Fibre Channel architecture provides various communication protocols. Each interconnected storage system within the architecture is referred to as a *node*, and each host is also a node. Each node corresponds to one or more ports. (In the case of Fibre Channel I/O adapters, these ports are Fibre Channel ports.) Each port attaches to a serial-transmission medium that provides duplex communication with the node at the other end of the medium. You can configure your network structure that is based on one of three basic interconnection topologies (network structures):
 - Point-to-point and switched fabric
 - Fibre Channel arbitrated loop
 - Fibre Channel (FICON). This connection does not apply to open systems.

See the *IBM DS8000 Host Systems Attachment Guide* for more information about these supported topologies.

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

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

Note: Do not use link extenders with emulation functions on links over which Remote Mirror and Copy operations are performed. Link extenders with emulation functions introduce more path delay.

- Because the Fibre Channel architecture allows any channel initiator to access any Fibre Channel device, without access restrictions, a security exposure can occur. Have your IBM service representative set the Fibre Channel access modes to the correct setting. See the *IBM DS8000 Host Systems Attachment Guide* for more information about Fibre Channel access modes.
- Storage systems can connect to IBM SAN Volume Controller host systems. See the *IBM DS8000 Host Systems Attachment Guide* for more information.

Attention: Signal-loss tolerance for Fibre Channel links that run at 8 or 16 Gbps are reduced by -1.4 dB as compared to links that run at 4 Gbps. Take necessary steps to ensure that links are within the signal-loss parameters that are listed when you plan a move from 4 Gbps FC to 8 or 16 Gbps FC adapters with existing infrastructure. Using more than two patch panels or other link connectors between two 8 or 16 Gbps ports at maximum distance of 10 km might result in greater link loss than is acceptable. Signal loss tolerance with measurements, including all connectors and patch panel connections, are:

- Loss per 10 km for 16 Gbps speed is -6.4 dB.
- Loss per 10 km for 8 Gbps speed is -6.4 dB.
- Loss per 10 km for 4 Gbps speed is -7.8 dB.

OM3 Fibre Channel cables are required to support 8 Gbps host adapters.

Chapter 7. Planning your storage complex setup

During installation, IBM customizes the setup of your storage complex that is based on information that you provide in the customization worksheets.

Each time that you install a new storage system or management console, you must complete the customization worksheets before installation by IBM service representatives.

Use the customization worksheets to specify the initial setup for the following items:

- Company information
- Management console network settings
- Remote support (includes call home and remote service settings)
- Notifications (includes SNMP trap and email notification settings)
- Power control
- Control Switch settings

Important: IBM service representatives cannot install a storage system or management console until you provide them with the completed customization worksheets.

Company information

Specify on the company information worksheet any information which IBM service personnel (or your service provider) can use to contact you as quickly as possible or to access your storage complex.

This information includes the following items:

- General company information, such as company name and telephone number
- Administrator contact information
- Storage complex location and modem numbers

You must complete this worksheet for all installations that include a management console.

Management console network settings

Use the management console network setting worksheet to specify the IP address and LAN settings for your management console.

The management console network settings include the following items:

- Management console network identification
- Ethernet settings, if you want the management console to connect to your LAN
- DNS settings, if you plan to use a domain name server to resolve network names
- Routings, if you want to specify a default gateway for routing

Note: IBM attaches your LAN after the storage complex is installed and in operation.

You must complete the worksheet for all installations that include a management console. Before you complete the worksheet, review the exceptions that are listed in the notes at the bottom of the worksheet.

Remote support settings

The remote support worksheets specify whether you want outbound (call home) or inbound (remote services) remote support, or both.

Ensure that you enable both outbound and inbound support to help you maintain the highest availability of your data.

When you enable outbound (call home) remote support, your management console sends an electronic call home record to IBM support when there is a problem within the storage complex. If inbound remote service is also enabled, an IBM service representative can securely sign on to the management console from a remote location in response to the service call.

You can choose Internet SSL (preferred), Internet VPN, or modem as connectivity options for outbound call home.

IBM offers Assist On-Site (AOS) as a secure inbound remote service option. AOS provides a mechanism to establish a secure network connection to IBM over the internet by using SSL encryption. This is different from the screen sharing technology that is also part of AOS. The management console can make a direct AOS connection to IBM if your security guidelines allow this type of configuration.

Note: For easiest network setup, configuration and usability IBM recommends Internet SSL and AOS for outbound and inbound connectivity. You can use these connectivity options to operate your storage system without a modem attached. IBM might discontinue VPN and modem as remote support option in future versions of the DS8000 product.

The management console can also be configured to offload error logs to IBM over the Internet by using FTP. This option is normally only used when the modem is configured for outbound remote support.

For any of the remote support connectivity methods, you can use the data storage command-line interface (DS CLI) and its audit log feature to review who performed any remote service on your storage system, and at what time the work was performed. Contact your IBM service representative for more information on which service actions are completed remotely. You can also use DS CLI to control network and remote service access to each management console and the storage system.

For additional information on SSL and VPN based remote support connectivity, go to <http://www-01.ibm.com/support/docview.wss?uid=ssg1S1002693>. The IBM AOS redbook at <http://www.redbooks.ibm.com/abstracts/redp4889.html?Open> provides additional information on AOS as a secure remote service solution.

You must complete the worksheet for all installations that include a management console.

Notification settings

Use the notification worksheets to specify the types of notifications that you want to receive and that you want others to receive.

Note: The IBM service representative sets up the notification process.

Notifications contain information about your storage complex, such as serviceable events.

You can receive notifications through the following methods:

- Simple Network Management Protocol (SNMP) traps
- Email

You can choose one or both notification methods.

When you choose to have your storage complex generate SNMP traps, you can monitor the storage complex over your network. You can control whether management information base (MIB) information is accessible and what type of SNMP traps to send. You can also specify the maximum number of traps that are sent for each event and where to send the traps.

Notes:

1. If you have open-systems hosts and remote mirror and copy functions, you must enable SNMP notifications for status reporting.
2. If you plan to use advanced functions SNMP messaging, you must set those functions by using DS CLI.

When you choose to enable email notifications, email messages are sent to all the email addresses that you specify on the worksheet when the storage complex encounters a serviceable event or must alert you to other information.

You must complete the worksheet for each management console to be installed.

Power control settings

If you use the remote z Systems power control feature, you must specify that on the power control worksheet, so the power mode can be set up to support that feature. Use this worksheet to specify and schedule whether power turns on and off automatically.

If you want to use a scheduled power mode, you must enter the schedule on the power control worksheet. You must complete the power control worksheet for all installations.

Control switch settings

Use the control switch settings work sheet to specify certain settings that affect host connectivity. You are asked to enter these choices on the control switch settings work sheet so that your IBM service representative can set them during the installation of your storage system.

The following control switches are set using the choices you specify on the control settings work sheet.

IBM i LUN serial suffix number

Use this control switch setting only when you attach more than one DS8000 storage system to an AS/400 or IBM i host *and* the last three digits of the worldwide node name (WWNN) are the same on any of the storage systems.

Control switch settings - attachment to IBM z Systems

The following control switch settings are specific to IBM z Systems.

16 Gb FC Forward Error Correction

Forward Error Correction (FEC) intercepts and corrects certain errors to improve the adoption of 16 Gbps FCP and FICON host adapters.

Control-unit initiated reconfiguration (CUIR) support

Control-unit initiated reconfiguration (CUIR) allows automation of channel path quiesce and resume actions during certain service actions. This setting eliminates the requirement for manual actions from the host.

Control unit threshold

This control switch provides the threshold level, presenting a SIM to the operator console for controller related errors. SIMs are always sent to the attached IBM z Systems hosts for logging to the Error Recording Data Set (ERDS). SIMs can be selectively reported to the IBM z Systems host operator console, as determined by SIM type and SIM severity.

Device threshold

This control switch provides the threshold level, presenting a SIM to the operator console for device related errors. Device threshold levels are the same type and severity as control unit threshold settings.

Lights-on fast load for the 8 Gb/s host adapters

8 Gb Fibre Channel lights-on fast load provides seamless microcode update to storage arrays, and eliminates potential loss of access.

"Fast load" is a general term for quickly loading a new code version to a host adapter card on the DS8000. "Lights-on fast load" is an enhanced feature of the fast load process. Before the introduction of this enhanced feature, the laser light that sends the Fibre Channel signals would be turned off briefly during a fast load. This interruption to the signals caused the host at the other end to go through time-consuming recovery actions. The lights-on fast load feature keeps the laser light turned on throughout the fast load process. There are still intervals when the signals are paused, but lights-on fast load keeps the light path active so that there is less disruption to the host when a new code version is loaded to the adapter card.

Lights-on fast load for the 16 Gb/s host adapters

16 Gb Fibre Channel lights-on fast load provides seamless microcode update to storage arrays, and eliminates potential loss of access.

Media threshold

This control switch provides the threshold level, presenting a SIM to the operator console for media related errors. Media threshold levels are the same type and severity as control unit threshold settings.

Present SIM data to all hosts

Service Information Messages (SIMs) are offloaded to the first I/O request directed to each logical subsystem in the storage facility if the request is device or control unit related, or offloaded to the individual logical volume

| when the request is media related. This control switch determines whether
| SIMs are sent to all, or to only the first, attached IBM z Systems LPAR
| making an I/O request to the logical system or logical volume.

| **IBM z Systems high-performance FICON enhanced buffer management**

| IBM z Systems high-performance FICON enhanced buffer management
| provides improved performance for multi-site configurations when writing
| data remotely (remote site recovery).

Chapter 8. Planning data migration

The planning and methods of data migration for the DS8000 series varies by environment.

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

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

Data

- How much data is to be migrated?

Operating system

- Is it a z Systems or UNIX system? Consider using IBM Remote Mirror and Copy functions such as Metro Mirror, Global Mirror, or some variation of a logical volume manager.
- Is it z/OS? Consider using DFDSS, though there are many choices.
- Is it VM? Consider using DASD Dump Restore or PTAPE.
- Is it VSE? Consider using the VSE **fastcopy** or **ditto** commands.

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

Storage system

- Are the storage systems involved the same type with the same level of licensed machine code?
- Are the storage systems different? If the storage systems are different, ensure that the new configuration is large enough to accommodate the existing data. You also want to ensure that the virtual disks are similar in configuration to the disk drives that they are replacing.

Time and complexity

- What duration of service outage can be tolerated? Typically data migration requires that updates or changes cease while the movement occurs. Also, depending on the amount of data that you are moving and your migrating method, data might be unavailable for an extended time, even several hours.
- Does the complexity and time that is involved require the services of IBM through International Global Services? Contact your IBM service representative for more information.

When you replace existing storage, partition the storage so that the virtual disks are similar in configuration to the disk drives that they are replacing. New configurations must be large enough to accommodate the existing data.

You might want to take advantage of this opportunity to do some remapping. The allocation and distribution of data are not required to be a straight one-to-one relationship, although that is possible. For instance, you can take advantage of using a maximum of 255 logical subsystems whereas the prior limitation was 32 logical subsystems.

Consider creating any new fixed block (FB) volumes with T10 DIF protection. This protection can be used on volumes to which data is migrated, even if the current host server volumes are not T10-protected. T10 DIF-protected volumes can be used even if the host server does not currently support T10 DIF.

Selecting a data migration method

The data migration method that you select must provide the best compromise between efficiency and impact on the system users. The selected method provides a simple but robust method that minimizes user impact.

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

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

Consider the following items to determine the best method for your data migration:

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

Notes:

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

Table 84 compares data migration options and lists advantages and disadvantages of each.

Table 84. Comparison of data migration options

Type	Example	Advantages	Disadvantages
OS / LVM Mirroring	Logical Volume Managers, (LVM) Veritas Volume Manager (VxVM), Windows Disk Administrator	Little or no application service disruption	Potential application delays

Table 84. Comparison of data migration options (continued)

Type	Example	Advantages	Disadvantages
UNIX or Windows Commands	cpio, cplv, dd, tar, backup restore; copy, sscopy, xcopy, drag and drop	Common, easy to use, tested	Length of service interruption varies; scripting prone to errors and more testing
Remote Copy	Synchronous Mirror (Metro Mirror); Asynchronous Mirroring (Global Mirror and Global Copy)	Operating system independent	Like storage device types needed
Third-party software packages	Data Migration (XoSoft); Backup /Restore (Tivoli, Legato, Veritas)	Some have little application service interruption, standard utilities	Cost of software; some have high application service interruption
Third-party migration appliances	IBM SAN Volume Controller, DataCore SANsymphony	Multiple heterogeneous storage vendors supported; migration cycles that are offloaded to appliance	Cost of migration appliance / service, application disruption to install / remove appliance

Chapter 9. Planning for security

The DS8000 series provides functions to manage data secrecy and networking security, including data encryption. Version 7.2 introduces more functions to enable the storage system for NIST SP 800-131A requirements.

Planning for data encryption

The DS8000 series support data encryption by using IBM Full Disk Encryption (FDE) and key servers, such as IBM Security Key Lifecycle Manager or Tivoli Key Lifecycle Manager.

To enable disk encryption on FDE drives, the storage system must be configured to communicate with two or more key servers, such as IBM Security Key Lifecycle Manager or Tivoli Key Lifecycle Manager. The physical connection between the Hardware Management Console and the key server is through a Internet Protocol network.

Planning for encryption is a customer responsibility. There are three major planning components to the implementation of an encryption environment. Review all planning requirements and include them in the installation considerations.

Planning for encryption-key servers

DS8870 requires at least two encryption-key servers and associated software for each site that has one or more encryption-enabled storage systems.

One encryption-key server must be isolated. An isolated encryption-key server is a set of dedicated server resources that run only the encryption-key lifecycle manager application and its associated software stack. This server directly attached to dedicated non-encrypting storage resources containing only key server code and data objects.

The remaining key servers can be of any supported key-server configuration. Any site that operates independently of other sites must have key servers for the encryption-enabled storage systems at that site.

DS8000 encryption environments are recommended to configure an external management console for high availability (feature code 1130).

Important: You are responsible for replicating key labels and their associated key material across all key servers that are attached to the encryption-enabled DS8000 system before you configure that key label on the DS8000 system.

You can configure each encryption-enabled storage system with two independent key labels. This capability allows the use of two independent key-servers when one or both key-servers are using secure-key mode keystores. The isolated key-server can be used with a second key-server that is operating with a secure-key mode keystore.

If you need support for a dual-platform key server on a storage system, the installation of Tivoli Key Lifecycle Manager interim fix 2 (V1.0.0.2 or later) is recommended to show both key labels in the GUI. If you intend to replicate keys between separate IBM z Systems Sysplexes by using ICSF with the

JCECCARACFKS keystore in secure-key mode and with the secure-key configuration flag set in Tivoli Key Lifecycle Manager, Tivoli Key Lifecycle Manager Fix Pack 3 (V1.0.0.3 or later) is required.

Note: Starting from v2.5, Tivoli Key Lifecycle Manager changed its name to IBM Security Key Lifecycle Manager. This is not to be confused with IBM Security Key Lifecycle Manager for zOS.

Planning for key lifecycle managers

DS8000 series support various key lifecycle managers.

If NIST 800-131A security conformance is required on your DS8870 with LMC R7.2, select the version of IBM Security Key Lifecycle Manager that is appropriate for your encryption key server host and connection network protocol requirements.

- If your encryption key server runs on an open system host and you do not plan to use the Transport Layer Security (TLS) 1.2 protocol with this key server, use IBM Security Key Lifecycle Manager V2.0.1 or later.
- If your encryption key server runs on an open system host and you plan to use the TLS 1.2 protocol with this key server, use IBM Security Key Lifecycle Manager V2.5 or later.
- If your encryption key server runs on a IBM z Systems host LPAR with z/OS and you do not plan to use the TLS 1.2 protocol with this key server, use IBM Security Key Lifecycle Manager for zOS V1.1.0.2 or later.
- If your encryption key server runs on a IBM z Systems host LPAR with z/OS and you plan to use the TLS 1.2 protocol with this key server, use IBM Security Key Lifecycle Manager for zOS V1.1.0.3 or later.

If NIST 800-131A security conformance is not required on your storage system, select the appropriate encryption key manager for your encryption key server host.

- If your encryption key server runs on an open system host, install IBM Tivoli Key Lifecycle Manager v1.0 or later.
- If your encryption key server runs on an IBM z Systems host LPAR with z/OS, install IBM Security Key Lifecycle Manager for zOS v1.0.1 or later.

The IBM Security Key Lifecycle Manager (SKLM) is an encryption key appliance for encrypting hardware devices, such as disk or tape. The SKLM appliance (model 961 feature code 1760, or model AP1 feature code 1761) is required when you use SKLM software with the DS8000 series. The model AP1 appliance can be used with DS8000 disk or with TS series tape products.

Note: You must acquire an IBM Security Key Lifecycle Manager license for use of the SKLM software that is ordered separately from the stand-alone server hardware. The SKLM license includes both an installation license for the SKLM management software and a license for encrypting drives.

Storage systems that are ordered with this feature code have a Linux operating system and Tivoli Key Lifecycle Manager software preinstalled. If you acquire the Tivoli Key Lifecycle Manager Isolated Key Server on your own, refer to the Tivoli Key Lifecycle Manager Installation and Configuration Guide (SC27-2741) for hardware and operating system requirements.

Security Key Lifecycle Manager for z/OS generates encryption keys and manages their transfer to and from devices in an IBM z Systems environment.

Planning for full-disk encryption activation

Full-disk-encryption drives are standard on the DS8000 series. These drives encrypt and decrypt at interface speeds, with no impact on performance.

Full disk encryption offerings must be activated before use, as part of the system installation and configuration. This installation and activation review is performed by the IBM Systems Lab Services team. To submit a request or inquiry, see the Storage Services website (www-03.ibm.com/systems/services/labservices/platforms/labservices_storage.html), and click **Contact now**.

You are responsible for downloading or obtaining from IBM, and installing designated machine code (such as microcode, basic input/output system code (called BIOS), utility programs, device drivers, and diagnostics that are delivered with an IBM system) and other software updates in a timely manner from the ibm.com website (www.ibm.com) or from other electronic media, and following the instructions that IBM provides. You can request IBM to install machine code changes; however, you might be charged for that service.

Planning for user accounts and passwords

Planning for administrative user and service accounts and passwords ensures that you use the best security practices.

Managing secure user accounts

Follow these recommended practices for managing secure user accounts.

Procedure

Complete the following steps to achieve the level of secure access for users that is required for your storage system.

1. Assign two or more storage administrators and two or more security administrators to manage your storage system. To preserve the dual control that is recommended for recovery key management, do not assign both storage administrator and security administrator roles to the same user. Change the password for both the default storage administrator and default security administrator user accounts, or delete the default user account after user accounts for other administrators are created.
2. Create one user account for each user who is authorized to access your storage system. Do not share a single user account between multiple users.
3. Assign appropriate user roles and scopes to user accounts in accordance with the storage management responsibilities of the user.
4. Review configurable user ID policies, and set the policies in accordance with your security objectives. The default settings are consistent with IBM recommended user ID and password policies and practices and are enabled without additional configuration.
5. For applications that require network access to the storage system, assign a unique user ID (in other words, an ID that is not assigned to any other user). You can assign different user IDs for different software applications or different servers so that actions can be distinguished by user ID in the audit logs.

Managing secure service accounts

Follow these recommended practices to manage access to your service account in the DS Service GUI and remote access by IBM Hardware Support.

Procedure

Complete the following steps to achieve the level of secure access that is required for service accounts on your storage system.

1. Assign one or more service administrators to manage service on your storage system.
2. Access the DS Service GUI from a web browser on a system that has network access to the Hardware Management Console (HMC) at `https://HMC_IP/service`, where *HMC_IP* is the IP address or host name of the HMC. You can also access the DS Service GUI from the link on the login page of the DS8000 Storage Management GUI.
3. Log in to the DS Service GUI by using the service administrator account and change the password for that account.

The service administrator account is pre-configured with user ID (customer) and password (cust0mer).

4. Determine how you want IBM Hardware Support to access your storage system and set remote service access controls accordingly.

Before installation of the storage system, your IBM service representative consults with you about the types of remote service access available. IBM recommends Assist On-site (AOS) as a secure remote service method. AOS provides a mechanism to establish a secure network connection to IBM over the internet with SSL encryption. It can be configured so that the service administrator must approve remote service access and can monitor remote service activity.

Planning for NIST SP 800-131A security conformance

The National Institute of Standards and Technology (NIST) SP 800-131A is a United States standard that provides guidance for protecting data by using cryptographic algorithms that have key strengths of 112 bits.

NIST SP 800-131A defines which cryptographic algorithms are valid and which cryptographic algorithm parameter values are required to achieve a specific security strength in a specific time period. Starting in 2014, a minimum security strength of 112 bits is required when new data is processed or created. Existing data processed with a security strength of 80 bits should remain secure until around 2031, subject to additional NIST standards with guidelines for managing secure data.

In general, DS8870 with licensed machine code (LMC) V7.2 allows the use of 112-bit security strengths if the other unit that is attached to the network connection supports 112-bit security strength. If security levels are set to conform with NIST SP 800-131A guidelines, the DS8870 requires 112-bit security strength on all SSL/TLS connections, other than remote service network connections.

On network connections that use SSL/TLS protocols, 112-bit security has the following requirements:

- The client and server must negotiate the use of TLS 1.2.
- The client and server must negotiate an approved cipher suite that uses cryptographic algorithms with at least 112-bit security strength.
- The client or server must limit hash and signature algorithms to provide at least 112-bit security strength; for example, the client must prevent the use of SHA-1 hashes.

- Certificates that are used by the client or server must have public keys and digital signatures with at least 112-bit security strength, such as RSA-2048 keys with SHA-256 digital signatures.
- Deterministic random bit generators (DRBGs) must use approved algorithms with a least 112-bit security strength and must be provided with entropy sources that have at least 112 bits of entropy.

To enable NIST SP 800-131A security conformance in your environment, update the following entities. It might not be feasible to update all of these entities at the same time because of various dependencies. Therefore, you can upgrade them for NIST SP 800-131A security conformance independently of each other.

- Encryption key servers
- Remote authentication servers
- DS Network Interface clients
- DS Network Interface server
- DS8000 Storage Management GUI and DS Service GUI servers
- SMI-S agents

Attention: Before you disable earlier SSL/TLS protocols on the DS8870, you must ensure that all external system networks connected to the DS8870 are enabled for TLS 1.2 and are NIST SP 800-131A compliant. Otherwise, network connection to these systems will be prohibited.

For information about configuring your environment for NIST SP 800-131A conformance, see security best practices in the IBM DS8000 series online product documentation (www.ibm.com/support/knowledgecenter/ST8NCA/product_welcome/ds8000_kcwelcome.html).

Chapter 10. License activation and management

The management and activation of licensed functions are responsibilities that are associated with the role of your storage administrator.

Management refers to the use of the IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa) to select a license scope and to assign a license value. You can complete these activities and then activate the function.

Note: If you are activating features for any of the licensed functions, such as Copy Services, all features must have the same capacity, including the operating environment license feature.

Activation refers to the retrieval and installation of the feature activation code into the storage system. The feature activation code is obtained by using the IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa) and is based on the license scope and license value.

You complete these activities at the following times:

- After the IBM service representative installs your storage system, and before you configure it
- When you increase the extent of the function authorization for a licensed function (that is, you add more capacity to your license)

To complete these activities, access the IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa).

When you access DSFA, you must enter information about your storage system so the web application can access the correct function authorization records. You can find this information either by clicking **Need Help** in the Activate Licensed Functions window or by selecting **Properties** on the System page of the DS8000 Storage Management GUI.

Planning your licensed functions

As you plan for your licensed functions, it is important to consider increases in your workload requirements. To provide more flexibility with licensing, use-based licensing is supported on some models in the DS8000 series.

DS8870 functions, such as Parallel Access Volumes (PAVs) and advanced Copy Services functions, FlashCopy, Metro Mirror, and Global Mirror are priced based on the use of license capacity that you purchase. For example, when you purchase licensed features for these functions, the features no longer must be purchased with the same license capacity. You can purchase licenses for Copy Services functions with individual capacities instead of having to buy the licensed features with a single capacity. This capability is designed to provide you with more flexible and optimal price and performance configurations.

With the use-based license capability comes the requirement to plan how much storage capacity you require for future growth. As you plan for your licensed functions, it is important to consider increases in your workload requirements. For example, consider the following guidelines, which include but are not limited to:

- Plan for storage space allocation. Determine your typical storage requirements and consider how much more storage you would need if you have rapid or unpredictable growth.
- Estimate the amount of capacity you need for current and future Copy Services functions. For example, consider the number of target volumes you need for FlashCopy relationships at remote sites. As the number of FlashCopy target volumes increase, more available bandwidth of the disk system might be used by the copy process. In addition, Copy Service solutions that require multiple copies of data can also require extensive capacity storage.

Recognizing that both your storage and data requirements will increase over time and that capacity and performance planning is your responsibility, ensure that you purchase and manage your licensed functions for maximum usage. It can be more cost effective to purchase more storage capacity to ensure that the maximum usage of your licensed features does not exceed the allowed capacity of the storage that was purchased. If the capacity is exceeded, IBM is notified that the usage exceeds the allowed capacity on any given licensed feature. You are notified by IBM and required to extend enablement of your licensed feature and install a new licensed feature key.

Activation of licensed functions

After the IBM service representatives complete your DS8000 series installation, your first step is to activate your licensed functions.

To activate your licensed functions, complete the following actions.

- Obtain your feature activation codes.
- Apply the activation codes to your storage system. You can apply the activation codes by importing a downloadable file from the IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa).

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

Note: The following activities are non-disruptive, but take effect at the next IML.

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

Activating licensed functions

You must obtain feature activation codes for the licensed features for each storage system by connecting to the IBM Disk Storage Feature Activation (DSFA) website.

Before you begin

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

- The IBM license feature authorization documents. If you are activating codes for a new storage system, these documents are included in the shipment. If you are activating codes for an existing storage system, IBM sends these documents to you in an envelope.

- A removable media for downloading your activation codes into a file. Use the removable media if you cannot access the Storage Manager from the system that you are using to access the DSFA website. Instead of using removable media, you can also write down the activation codes and then manually enter them into the system that runs the Storage Manager.
- The system serial number, model, and signature.

Notes:

1. Enabling an optional licensed function is a concurrent activity (assuming the appropriate level of microcode is installed for the function). The following activating activities are non-disruptive, but take effect at the next IML:
 - Removal of a licensed function for its deactivation.
 - A lateral change or reduction in the license scope. A lateral change is defined as changing the license scope from fixed block (FB) to count key data (CKD) or from CKD to FB. A reduction is defined as changing the license scope from all physical capacity (ALL) to only FB or only CKD capacity.
2. Before you begin this task, you must resolve any current system problems. Contact IBM Hardware Support for assistance.
3. Before you configure, disable or provide paths through any firewall because it might interfere with system communication.

About this task

Complete the following steps to activate license codes.

Procedure

You can activate the licensed functions from one of two locations in the DS8000 Storage Management GUI: the System Setup wizard during initial configuration; or the **Licensed Functions** tab of the System settings page.

1. Click **Activate Licensed Functions** or **Activate**.
2. Enter the license keys.
 - If you received your license keys from an IBM service representative, enter them in the Activate Licensed Functions window.
 - If you need to obtain your license keys from the IBM Data storage feature activation (DSFA) website, complete the following steps.
 - a. Go to IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa).
 - b. Click **DS8000 series**.
 - c. Enter the machine type, serial number, and machine signature of your DS8000 system. You can find this information either by clicking **Need Help** in the Activate Licensed Functions window or by selecting the **Properties** action on the **System** page.
 - d. Download the license keys XML file.
 - e. In the Activate Licensed Functions window, click the **Browse** icon to select the license keys XML file that you downloaded and click **Open**.
3. Click **Activate**.

Scenarios for managing licensing

These topics provide scenarios for managing your licenses after you initially activate them.

The following scenarios are provided:

- Adding storage capacity to an existing storage system
- Managing an active license feature

Note: More scenarios can be found on the IBM DS8000 Information Center.

Adding storage to your machine

You can add storage (in terabytes) to an existing licensed feature, such as point-in-time copy.

About this task

For this scenario, assume that you initially purchased 23 TB of point-in-time capacity. You can order two 7253 features (10 TB each) and three 7251 features (1 TB each). After several months, you need an extra 20 TB for your point-in-time copy operations. To increase storage, you must purchase and activate a larger license. This activity is nondisruptive and does not require that you restart your storage system.

Procedure

1. You order two of feature 7253 (10 TB each) against the serial number of the 242x machine type license currently on your storage system. These features are the additional features that increase your point-in-time copy authorization level.
2. After you order the features, you receive confirmation from IBM that these new features were processed.
3. Connect to the IBM-supported Disk Storage Feature Activation (DSFA) website at IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa) to retrieve an activation code for the licensed feature. This new code represents the total capacity that is now licensed (or 45 TB). It licenses the original 25 TB plus the additional 20 TB that you just ordered.
4. After you obtain the activation codes for the licensed feature, enter it into the web-based DS8000 Storage Management GUI. You replace the existing activation code with the new activation code.
5. After the activation code is installed successfully, you now have 45 TB of PTC capacity.

Managing a licensed feature

Use the IBM Data storage feature activation (DSFA) website to change an optional function from active to inactive. Change an assigned value, such as current number of terabytes, for a feature to make that licensed feature inactive.

About this task

If you have an active optional function and you want to replace it with an inactive feature, you must repurchase the feature if you want to use it again. However, you can use the following steps if you want to use the feature again.

Procedure

1. From the IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa), change the assigned value from the current number of terabytes (TB) to 0 TB.
2. If this change is made, you can go back to the DSFA website and reactivate the feature, up to the previously purchased level, without having to repurchase the feature.

Appendix A. Accessibility features for IBM DS8000

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

These are the major accessibility features associated with the IBM DS8000 series online product documentation.

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. HTML documents have been tested using JAWS version 15.0.
- This product uses standard Windows navigation keys.
- Interfaces that are commonly used by screen readers.
- Keys are discernible by touch but do not activate just by touching them.
- Industry-standard devices, ports, and connectors.
- The attachment of alternative input and output devices.

The DS8000 online product documentation and its related publications are accessibility-enabled. The accessibility features of the online documentation are described in the IBM Knowledge Center website (www.ibm.com/support/knowledgecenter/).

Keyboard navigation

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 DS8000 online documentation from the keyboard by using the shortcut keys for your browser or screen-reader software. See your browser or screen-reader software Help for a list of shortcut keys that it supports.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (www.ibm.com/able/) for more information about the commitment that IBM has to accessibility.

Appendix B. Warranty information

The statement of limited warranty specifies the warranty period, type of warranty service, and service level.

See IBM Warranty Information for more information.

Appendix C. IBM DS8000 equipment and documents

Use the documents that IBM provides with your DS8000 series to identify and check your main components.

The equipment that you receive can be grouped as follows:

- Components that must stay with the shipment because they are needed for installation
- Components that are for customer use
- Components that must stay with the storage system after installation because they are needed by IBM service representatives

Note: These lists are not intended to be a comprehensive list. They describe only the main shipped components.

Installation components

Your shipment includes all the equipment that is needed for the installation of your storage systems. Equipment includes storage systems, power cords, adapters, cables, installation instructions, and other essential material.

The following installation components are included with your shipment:

Storage system

Your shipment includes one or more of the following frames that you ordered:

- Base frame (model 961)
- Expansion frames (model 96E)

When the frames arrive, they contain any ordered I/O enclosures, device adapters, storage enclosures, drives, and the appropriate cables to support those components. IBM installs these components at the factory.

Hardware Management Console

A laptop-based Management Console is included with each base frame that you order. The management console is physically located (installed) inside the base frame.

If you order an external management console, you receive this separate workstation, the necessary Ethernet cables, and any required installation hardware to install the external management console in a rack that you provide. The IBM service representatives install the external management console into an IBM or a non-IBM rack when they install your storage system.

Power cords

Your shipment includes the country or region-specific power cord that you ordered.

Various media

IBM ships the following media (typically CDs), which are used during the installation of your storage systems:

- Installation media, which includes installation scripts for the I/O attachment for AIX and HP-UX, IBM DS CLI (command-line interface) software, and IBM Multipath Subsystem Device Driver installation instructions and software
- Licensed machine code (LMC) media for the MC
- Operating system media
- LMC media for the /242x machine type
- Quick Code Reference document that details program code, utilities, and documentation included in the ship group

Hardcopy installation instructions

Your shipment includes hardcopy installation instructions for the IBM service representatives who install your storage system.

Engineering changes (if applicable)

IBM occasionally releases engineering changes (ECs) to correct problems or provide more support. If released, these ECs are included in your shipment for the IBM service representative to install.

Customer components

IBM provides media and documents that are intended for you to keep.

- License and warranty documents
- READ ME FIRST for IBM products
- Quick Code Reference, which includes a listing of customer publications and media that is provided with the storage system
- Customer publications CDs: One CD contains PDFs of customer publications and the other CD contains PDFs of license and warranty documents.

Service components

IBM provides service-related media and documents with your storage system.

Keep the following components with your storage system so that IBM service representatives can use them when they service your storage system.

Service media

Your delivery includes the following media for IBM service representatives to use:

- Operating system media
- Management console media:
 - Management console critical backup SDHC memory card
 - Dump, trace, statesave SDHC memory card, which IBM service representatives use for extracting statesave information during service
- A program temporary fix (PTF) CD for the operating system
- Service documents CD, which includes the following documentation: DS8000 service documentation and the DS8000 parts catalog.

Appendix D. DS8800 to DS8870 model conversion

You can convert an existing DS8800 (models 951 and 95E) into an Enterprise Class, high-performance DS8870 (model 961 and 96E).

DS8800 to DS8870 model conversion summary

The conversion is a non-concurrent operation, requiring your planning, several days of onsite IBM support and your full participation.

Model conversion summary:

- The model conversion does not change the machine type.
- Each DS8800 that is converted to DS8870 retains the existing (original DS8800) frame serial numbers, worldwide node name (WWNN), and worldwide port names (WWPNs).
- IBM provides new frames, including the management console, processor enclosures, power supplies, and I/O enclosures. The first expansion frame includes power supplies and I/O enclosures. The second and third expansion frames include power supplies.
- The device adapters, host adapters, storage enclosure assemblies, and disk drives from the existing DS8800 frames are moved into the new DS8870 frames.
- Each DS8800 frame or frames (without the adapters, storage enclosures and disk drives) is returned to IBM.
- Applicable licensed machine-code keys from DS8800 need to be reapplied to the new DS8870 after conversion and before logical configuration is created.
- Any additional upgrades that are planned for DS8870 must be done after the conversion is completed.
- The existing DS8800 warranty is applied to the new DS8870 without extensions.
- The new DS8870 has less net capacity on the same drives because about 32 GB per rank is required for formatting and metadata considerations. If your existing DS8800 is near capacity before the model conversion, plan to purchase more storage. Otherwise, the net capacity on the drives might not be sufficient.
- If the existing DS8800 has a secondary management-console feature that is associated with it, a DS8870-compatible secondary management console is included when the model conversion MES involves a DS8800 system that is managed by pair of management consoles (primary and secondary). The secondary management console (if installed on source DS8800) must have the same drive capacity and memory as the base DS8870.
- The DS8870 storage systems that are converted from DS8800 storage systems are either all non-FDE or all FDE at the time of conversion. However, with release 7.4 or later, high-performance flash enclosures can be added to a converted DS8870 non-FDE system.

You must complete all of the prerequisites in the following table before the IBM service representative starts the model conversion:

Table 85. Model conversion prerequisites

Prerequisites	Description
Data migration, logical configuration removal, encryption group removal	You need to migrate all data from DS8800, remove all logical configuration, and remove the encryption group (if encryption is activated) from DS8800.
Secure data overwrites	DS8800 frames are required to be returned to IBM, including the processors and management console. If you require that data is securely removed from the management console and processor disk drives, you can request the IBM service representative to perform a secure-data overwrite. This is an extra service that is performed by the IBM service representative. The secure-data overwrite must be run after you have migrated data and removed the logical configuration and encryption group, and before you begin the conversion. The SDO includes an overwrite on the management console, processor disk drives, and customer data drives. This process can take significant time to complete. You need to plan accordingly.
Data encryption and Security Key Lifecycle Manager servers	If DS8800 uses data encryption and Security Key Lifecycle Manager servers and you intend to activate encryption on DS8870, ensure that the infrastructure is in place for DS8870 system. Encryption is optional for conversion. If the converted DS8870 has FDE drives, and you choose to activate encryption, ensure that you apply the encryption LIC keys and the Security Key Lifecycle Manager infrastructure is in place.
Ethernet cable, Fibre Channel host cables, and telephone line	Order the following cables for the new DS8870 only if your existing DS8800 infrastructure cannot be used: <ul style="list-style-type: none"> • Fibre Channel host cables from the host to DS8870 host adapters. • An Ethernet LAN signal cable. This cable is required to connect the management console to your LAN. • A functional analog telephone line to the management console. The phone jack needs to be within 15.24 meters (50 feet) of the management console. <p>Note: If you want the IBM service representative to route the cables, that activity must be coded or billed separately from the DS8870 installation work.</p>
Power cables	DS8870 uses a new power system and requires a unique set of power cords..
Power ratings	IBM does not recommend the use of GFCI, ELCB, or RCCB for mainline power. However, some sites require the use of GFI.
Earthquake Resistance Kit	The DS8870 earthquake resistance kit requires that you install a new set of hold down hardware in the concrete floor. The existing holes and hardware that is used for DS8800 cannot be reused.

Table 85. Model conversion prerequisites (continued)

Prerequisites	Description
Physical installation requirements	<p>Determine whether the new DS8870 can occupy the same floor space as the existing DS8800.</p> <ul style="list-style-type: none"> • If you are using the same location, you need to provide a location for both the existing DS8800 and converted DS8870 frames so that the hardware can be staged between them. • Whether installed in an existing or new location, the location must conform to the requirements for DS8870.
Management console settings	<p>Determine how the IBM service representative receives the settings for the new management consoles (primary and secondary).</p> <ul style="list-style-type: none"> • If the existing settings are to be reused, the IBM service representative records the existing settings to use on the new management console. • If some or all of the settings are changing, you must complete the customer configuration worksheets and provide the completed worksheets to the IBM service representative before you begin the conversion.
Single management console (1x1) configuration	<p>DS8800 must be in a 1x1 configuration. It must not be connected to another storage system by the black and gray private network cables. If it is not in a 1x1 configuration, the IBM service representative needs to use discontinue procedures to separate it before you begin the conversion. If the existing DS8800 is in either a 2x2 or 2x1 configuration, a significant amount of extra time is required to reconfigure the storage system to a 1x1 configuration before you begin the conversion. You need to plan accordingly.</p>
Serviceable events	<p>Before you begin the conversion, the IBM service representative must repair any open serviceable events. Fixing any hardware repairs adds service time.</p>

Checking your preparations

Use the following steps as a checklist to ensure that you complete all of the necessary planning and preparations before you receive shipment of the storage system.

Procedure

1. Verify that you have two ac power sources for each frame.
2. Verify that the power source matches the power cord feature code for the storage system. The power connector type for the current model differs from the previous models.
3. Verify that the circuit breaker rating meets the storage system requirements.
4. Verify that the physical requirements necessary storage system installation are sufficient.

The location, floor type, frame weight, and cable management must be considered.

5. Verify that you have a complete list of the necessary licensed function keys. The list of licensed function keys must match that of the converted model. The keys are used to reactivate all licensed functions after completion of the model conversion.

Removing data, configuration, and encryption

You must migrate all data from the existing storage system, remove the logical configuration, and remove the encryption group before the IBM service representative can begin the model conversion.

Procedure

Ensure that the following operations are completed on the storage system that is to be converted before the IBM service representative arrives:

1. Migrate all data from the storage system.
2. Remove the logical configuration from the storage system.
3. Remove the encryption group, recovery keys, and Tivoli Key Lifecycle Manager server information from the storage system.
4. If data must be securely removed from the management console and processor drives, you can request the IBM service representative to perform a Secure Data Overwrite (SDO). This is an extra service that can be performed by the IBM service representative. The SDO must be performed after you have migrated data and removed the logical configuration and encryption group, and before you begin the conversion. An SDO can take significant time to complete. You must plan accordingly.

Completing post-conversion tasks

After the model conversion is completed by an IBM service representative, a few more tasks must be completed.

Procedure

Complete the following tasks on DS8870:

1. Download the necessary licensed-machine-code keys from the IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa) by using the same frame serial number as the storage system that was converted.
2. Apply the licensed-machine-code keys to enable the operating environment license, encryption (if applicable), and advanced functions.
3. If encryption is to be activated:
 - a. Assign the Security Key Lifecycle Manager servers.
 - b. Create the encryption recovery keys to prevent encryption deadlock.
 - c. Create the encryption groups.
4. Create the logical configuration. The configuration cannot be cloned from the converted storage system.

Appendix E. Customization worksheets

You must complete the customization worksheets before any installation of a new storage system or management console. After you fill out the worksheets, give them to the IBM service representatives who complete the installation.

Use the following worksheets to specify the initial settings for your storage system. You can customize settings for company information, the management console network, remote support settings, notifications, and power control.

Company information worksheet

Complete the company information worksheet to provide basic information about your company and administrator, and general system information.

Purpose

IBM service representatives use the information that is provided on the company information worksheet to customize your storage system. When you use any of the remote support features, the management console (MC) sends this information to IBM so an IBM service representative can contact you.

You must complete the Table 86 worksheet for all installations that include a management console.

Table 86. Company information worksheet

Item or setting	Instructions	Your information
Company name	Provide the name of your company. IBM service representatives use this information to identify your company when they receive call home reports from your storage system. Ensure that the company name that is provided here is consistent with all other systems that correspond to your IBM customer account.	
Customer number	Provide the customer number that is assigned by IBM to your company.	
Administrator information: Provide information about your storage system administrator in the following section.		
Administrator name	Provide the name of the individual at your site who service representatives can contact about storage system service matters.	
Administrator email address	Provide the email address that can be used to contact the administrator.	
Administrator telephone number	Provide the primary telephone number for service personnel to use to contact the storage system administrator. Include the area code and the country code, if appropriate.	

Table 86. Company information worksheet (continued)

Item or setting	Instructions	Your information
Alternate telephone number	Provide an alternative or off-shift telephone number that IBM service representatives can use to contact the storage system administrator. Include the area code and the country code, if appropriate.	
Fax number	Provide the primary fax number that IBM service representatives can use when they must fax documents to the storage system administrator. Include the area code and the country code, if appropriate.	
Alternate fax number	Provide an alternative fax number that service personnel can use when they must fax documents to the storage system administrator. Include the area code and the country code, if appropriate.	
Administrator's mailing address	Provide the mailing address for the administrator. Specify the complete address, including the street address, building (if appropriate), city or locality, state or province, and postal or zip code.	
Storage system information: Provide basic information about your storage system and the management console in the following section.		
System location	If different from the administrator's address, provide the full address where the storage system is located. Include the street address, building (if appropriate), city or locality, state or province, and postal or zip code.	
Modem number (MC1)	For the first MC to take advantage of inbound remote services, Provide the telephone number to the modem of the first MC in the storage complex. Include the area code and the country code, if appropriate. Note: Internet SSL / AOS is the preferred option for inbound remote services.	
Modem number (MC2, if installed)	For a second MC to take advantage of inbound remote services, provide the telephone number to the modem of the second MC. Include the area code and the country code, if appropriate. Note: Internet SSL / AOS is the preferred option for inbound remote services.	

Management console network settings worksheet

Specify basic network information on the management console network settings worksheet. Indicate your management console identification, Ethernet settings, DNS settings, and other required information.

When your storage system sends IBM any call home information through VPN or sends you notices about serviceable events (by using SNMP trap or email), these settings are included in the information to identify and provide important information about the management console that sent a service request.

Worksheet purpose

IBM service representatives use the information that you provide on the management console network settings worksheet to set up the network settings that support your management console.

You must complete this worksheet for all installations that include a management console.

Worksheet and instructions

The IP addresses and name server information that are shown on this worksheet are examples only. Contact your IBM service representative for the correct IP addresses and name servers that function with your storage system. Complete the Table 87 worksheet to provide the IBM service representatives information about how to set up your network settings.

Table 87. Management console network settings worksheet

Item or Setting	Instructions	MC1	MC2 (if applicable)
MC name (console name)	Record the console name that you want to assign to the management console workstation (for example, dsmc1). The console name and the domain are used to identify the MC to the network. Important: The name of each management console must be unique. Do not use the same name for the primary and secondary management consoles.		
Domain name	Provide the domain name that you are assigning to the MC workstation (for example, medina.xyz.it).		
Console description	Optional field to enter more details to further describe this storage facility.		

Table 87. Management console network settings worksheet (continued)

Item or Setting	Instructions	MC1	MC2 (if applicable)
MC time and time zone	Check Use local time and time zone to set the MC to the local time and time zone (local to where the MC is physically located). If you want to use a different time zone, check Use the following time zone and specify the time zone to use (for example, Central , European Time , or US Pacific Time).	_Use local time and time zone _Use the following time zone: _____	_Use local time and time zone _Use the following time zone: _____
NTP time server	List one or more external NTP time servers and the NTP protocol version. The time server must be reachable from the MC. Only MC1 needs to be set up for an external time server.	NTP time server IP address or host name: NTP protocol version: NTP time server IP address or host name: NTP protocol version:	Not applicable
Ethernet settings: Complete the LAN adapter details section when the MC connects to your LAN.			
Media speed (Ethernet)	Check Autodetection or the media speed of the Ethernet adapter. Tip: If you check Autodetection , the MC can automatically select the media speed appropriate for your configuration.	_Autodetection _10 Mbps Half Duplex _10 Mbps Full Duplex _100 Mbps Half Duplex _100 Mbps Full Duplex _1000 Mbps Half Duplex _1000 Mbps Full Duplex	_Autodetection _10 Mbps Half Duplex _10 Mbps Full Duplex _100 Mbps Half Duplex _100 Mbps Full Duplex _1000 Mbps Half Duplex _1000 Mbps Full Duplex
TCP/IP interface address	Ensure that the TCP/IP address that you select is not in a range that is reserved for the 242x private network. For more information, see "TCP/IP address ranges" on page 222. If you plan to use the IBM Internet VPN connectivity with the Network Address Translation (NAT), the IP address for the MC must be a routable private address (RFC1981) that is not already available on the Internet. Record the dotted decimal address that you are assigning to the MC (for example, 7.127.152.14).		

Table 87. Management console network settings worksheet (continued)

Item or Setting	Instructions	MC1	MC2 (if applicable)
242x private network address ranges	Can the default 242x private network address ranges be used? For more information, see “TCP/IP address ranges” on page 222.	<input type="checkbox"/> Yes: Use default ranges. <input type="checkbox"/> No: Check off the address range option to use. Table 89 on page 223 provides the addresses that are associated with each option. <input type="checkbox"/> Option 1 <input type="checkbox"/> Option 2 <input type="checkbox"/> Option 3	<input type="checkbox"/> Yes: Use default ranges. <input type="checkbox"/> No: Check off the address range option to use. Table 89 on page 223 provides the addresses that are associated with each option. <input type="checkbox"/> Option 1 <input type="checkbox"/> Option 2 <input type="checkbox"/> Option 3
IPv4 and IPv6 configuration	Select if the MC is connected to an IPv4 network, an IPv6 network, or a network that supports both IPv4 and IPv6. Select IPv6 only if you plan on assigning or autoconfiguring IPv6 addresses.	IPv4 only <input type="checkbox"/> IPv6 only <input type="checkbox"/> Both IPv4 and IPv6 <input type="checkbox"/>	IPv4 only <input type="checkbox"/> IPv6 only <input type="checkbox"/> Both IPv4 and IPv6 <input type="checkbox"/>
TCP/IP interface network mask	Record the dotted decimal network mask that you want to apply to the TCP/IP address (for example, 127.123.546.0).		
DNS settings: Complete the following section if you plan to use a domain name server (DNS) to resolve network names.			
Name server (DNS) internet address 1	Provide the dotted decimal address of the name server to be accessed by the MC (for example, 5.127.42.25).		
Name server domain name 1	Provide the domain name of the name server (for example, medina.xyz.it).		
Name server (DNS) internet address 2 (Optional)	Provide the dotted decimal address of the second name server that this workstation can access (for example, 5.127.42.252). Tip: You can specify a second name server when you configure a backup or secondary server for Copy Services.		
Name server domain name 2	If you have a second name server, provide the domain name of the second name server (for example, medina2.xyz.it).		
Routing settings: Complete the following section if you want to specify a default gateway for routing. (See Note following this table)			

Table 87. Management console network settings worksheet (continued)

Item or Setting	Instructions	MC1	MC2 (if applicable)
Gateway address	Confirm and record the dotted decimal or symbolic name address of the gateway (for example, 8.127.155.254 or londongate).		
<p>Notes:</p> <ol style="list-style-type: none"> Options that are in bold in the MC1 and MC2 columns indicate default settings. Management console is abbreviated as MC in the worksheet. 			

TCP/IP address ranges

When you select a management console TCP/IP address, ensure that you do not create a TCP/IP address conflict between your network and the 242x private network.

Use the following guidelines to prevent an address conflict:

- The management console TCP/IP address that you select must not be in the same address range as the address used by the 242x private networks.
- The TCP/IP addresses used by the 242x private networks must be outside the address ranges used by any network that the management console can reach.
- The 242x private network has one default address range. If the default address range cannot be used because it conflicts with another network, you can instead specify one of three optional addresses ranges.
- The IBM service documentation can refer to the two private networks as “black” and “gray” regardless of which address range was assigned.
- Table 88 on page 223 and Table 89 on page 223 can help you determine the 242x private network address range that the IBM service representative sets during the installation.

Use Table 88 on page 223 to determine whether the default address range can be used.

Table 88. Default TCP/IP address range determination

Question:	If the answer is no...	If the answer is yes...
<p>Do any of your networks to be reached by the management console use either of these address ranges?</p> <ul style="list-style-type: none"> • 172.16.0.0 to 172.16.255.255 (172.16.0.0/16) • 172.17.0.0 to 172.17.255.255 (172.17.0.0/16) 	<p>There is not a potential address conflict. The default 242x private network address ranges can be used. In Table 87 on page 219, use the following instructions:</p> <ol style="list-style-type: none"> 1. In the MC1 or MC2 column of the TCP/IP interface address row, as appropriate, enter the IP address for the management console. 2. In the MC1 or MC2 column of the 242x private network address ranges row, as appropriate, check Yes. 	<p>There might be an address conflict between your networks and the 242x private networks.</p> <p>Your service representative must reconfigure the 242x private networks to use one address range option from Table 89.</p> <ol style="list-style-type: none"> 1. In Table 89, select only one address range option for the 242x private network that does not create a potential address conflict. 2. In Table 87 on page 219, in the MC1 or MC2 column of the 242x private networks address ranges row, as appropriate, check the address range option that you selected from Table 89.

If the default address range cannot be used, use one of the optional TCP/IP address range options that are provided on Table 89.

Note: The options in the table are listed from most preferable (Option 1) to least preferable (Option 3).

Table 89. Address ranges of the storage facility private network (LIC bundles 75.xx.xx.xx and above)

Setting	Address range	
	Black network (MC eth0)	Gray network (MC eth3)
Option 1	10.235.0.0 to 10.235.2.255	10.236.0.0 to 10.236.2.255
Option 2	192.168.160.0 to 192.168.162.255	192.168.240.0 to 192.168.242.255
Option 3	9.15.0.0 to 9.15.2.255	9.16.0.0 to 9.16.2.255

IPv6 configuration

Use the options that are provided in Table 90 to configure the management console for communication over an IPv6 network.

Table 90. IPv6 configuration options

Item or setting	Instructions	MC1	MC2
IPv6 Autoconfig options	Select the autoconfig method for connecting the MC to an IPv6 network. You can check more than one option.	<p>Autoconfigure IP addresses []</p> <p>Use privacy extensions for autoconfiguration []</p> <p>Use DHCPv6 to configure IP settings []</p>	<p>Autoconfigure IP addresses []</p> <p>Use privacy extensions for autoconfiguration []</p> <p>Use DHCPv6 to configure IP settings []</p>

Table 90. IPv6 configuration options (continued)

Item or setting	Instructions	MC1	MC2
Static IPv6 address	List the static IPv6 addresses that you want to assign to the MC. Example IPv6 address: fe80:0:0:214:5eff:fe74:7ca8 Example prefix: 64	IPv6 address: Prefix length: IPv6 address Prefix length:	IPv6 address: Prefix length: IPv6 address Prefix length:

Remote support worksheets

The remote support worksheets specify the outbound (call home) and inbound (remote services) settings.

Use these worksheets to specify the settings to use for outbound and inbound remote support.

Important: Prerequisites, like firewall changes or analog modem phone lines, must be met at the time of the installation.

There are two remote support worksheets:

- Outbound (call home) worksheet
- Inbound (remote services) worksheet

Outbound (call home and dump/trace offload) worksheet

The outbound worksheet provides settings for the IBM call home feature and the dump/trace delivery method.

When the IBM service representative sets up the call home feature, the storage system automatically requests service from an IBM service representative when the system identifies a serviceable event. IBM service representatives can then analyze the problem without your having to alert them. The memory dump and trace offload transmissions provide the necessary information for IBM service representatives to quickly troubleshoot serviceable events.

Note: The IBM call home feature must be configured if you participate in the Standby Capacity on Demand program.

Worksheet purpose

IBM service representatives use the information that is provided on the outbound (call home) worksheet to customize your storage system to use or disable the call home feature.

Worksheet and instructions

Table 91 lists instructions for the outbound (call home) worksheet.

Table 91. Outbound (call home) worksheet

Item or Setting	Instructions	MC1	MC2 (if applicable)
Enable Call Home feature?	If you want IBM to enable the IBM call home feature, check Yes and then check the modes to use for call home. (You can select one or more modes. Internet SSL is the preferred option and it is selected by default. The attempted order of connectivity options is Internet SSL, Internet VPN, and then modem.)If you choose not to enable IBM call home, check No . If you select No , this worksheet is complete.	<input type="checkbox"/> Yes Modes: <input type="checkbox"/> By modem <input type="checkbox"/> By Internet SSL ¹ <input type="checkbox"/> By Internet VPN ² <input type="checkbox"/> No	<input type="checkbox"/> Yes Modes: <input type="checkbox"/> By modem <input type="checkbox"/> By Internet SSL ¹ <input type="checkbox"/> By Internet VPN ² <input type="checkbox"/> No

Table 91. Outbound (call home) worksheet (continued)

Item or Setting	Instructions	MC1	MC2 (if applicable)
<p>By Internet SSL is the preferred option for call home. By Internet VPN and the modem might be discontinued in a future version of this product.</p>			
<p>¹ If you select By Internet SSL, you allow your management console to use a secure sockets layer (SSL) connection over the Internet when a connection is established to the IBM service center.</p>			
<p>Internet SSL is the recommended option. If you select Internet SSL, it is recommended that you select AOS in the Inbound worksheet.</p>			
<p>You must open port 443:tcp in your network infrastructure for the following destination servers:</p>			
<p>Americas</p> <ul style="list-style-type: none"> • 129.42.160.48 • 129.42.160.49 • 207.25.252.200 • 207.25.252.204 			
<p>Non-Americas</p> <ul style="list-style-type: none"> • 129.42.160.48 • 129.42.160.50 • 207.25.252.200 • 207.25.252.205 			
<p>Problem Reporting Servers</p> <ul style="list-style-type: none"> • 129.42.26.224 • 129.42.34.224 • 129.42.42.224 			
<p>Configuration File Servers</p> <ul style="list-style-type: none"> • 129.42.56.216 • 129.42.58.216 • 129.42.60.216 			
<p>² If you select By Internet VPN, you allow your management console to use VPN over an Internet connection when a connection is established to the IBM service center. The IBM VPN implementation is a client/server VPN that is only active when it is needed. The two VPN endpoints are on the management console and on the IBM Boulder and the IBM Rochester VPN servers. There is no need for more VPN hardware in your network infrastructure.</p>			
<p>If you use VPN, the management console must have access through your Internet firewall to the following servers:</p> <ul style="list-style-type: none"> • IBM Boulder VPN server (IP address 207.25.252.196) • IBM Rochester VPN server (IP address 129.42.160.16) 			
<p>The first package is always sent from the management console. Only the following ports need to be open to the mentioned servers to use VPN:</p> <ul style="list-style-type: none"> • Port 500 UDP • Port 4500 UDP 			
<p>Your IBM service representative can also provide you with more information.</p>			
<p>If VPN or SSL connectivity are not options, you can configure FTP to offload log and trace data faster.</p>			
<p>For information about the IBM SSL and VPN remote support connectivity implementation, including technical details, see the website (www.ibm.com/support/docview.wss?uid=ssg1S1002693).</p>			
<p>Modem settings: Complete the following section if you are enabling call home by using a modem. Note that the modem might be discontinued in a future version of this product.</p>			
<p>Even if your VPN or SSL is configured, a modem can serve as a backup. If the modem is configured, IBM service representative can remotely access the management console without having a person onsite to initiate a connection.</p>			

Table 91. Outbound (call home) worksheet (continued)

Item or Setting	Instructions	MC1	MC2 (if applicable)
Dialing mode used by modem	Check the box that indicates whether your local modem uses tone or pulse dialing.	<input type="checkbox"/> Tone <input type="checkbox"/> Pulse	<input type="checkbox"/> Tone <input type="checkbox"/> Pulse
Wait for dial tone?	If the management console modem waits for a dial tone before you dial a telephone number, check Yes . If the modem dials without checking for the dial tone, check No .	<input type="checkbox"/> Yes <input type="checkbox"/> No	<input type="checkbox"/> Yes <input type="checkbox"/> No
Dial-out prefix	If the modem must dial a prefix to access an outside line, check Prefixes and provide the prefix numbers. If no prefix is necessary, check N/A .	<input type="checkbox"/> Prefixes: _____ <input type="checkbox"/> N/A	<input type="checkbox"/> Prefixes: _____ <input type="checkbox"/> N/A
Internet (SSL) settings: Complete the following section to select Internet (SSL) settings.			
Use SSL proxy?	If there is an SSL proxy server in between the HMC and the Internet, provide the address.	<input type="checkbox"/> Yes _____ <input type="checkbox"/> No	<input type="checkbox"/> Yes _____ <input type="checkbox"/> No
Authenticate with the SSL proxy?	If the SSL proxy requires a login, provide the user ID, and password.	<input type="checkbox"/> Yes _____ _____ <input type="checkbox"/> No	<input type="checkbox"/> Yes _____ _____ <input type="checkbox"/> No
Dumps and traces: Complete the following section to specify how you want to send trace or dump files to the IBM service center.			
Delivery method for dumps and traces	Check Use the call home setup if you want the management console to use the call home settings for dump and trace call-home transmissions. Check Use FTP if you want to send dumps and traces directly to the IBM FTP data repository server. FTP is recommended only if SSL or VPN connectivity is not provided. FTP is usually configured only if the modem is the only call home connectivity method. (To use FTP, you must connect the management console to your LAN and provide a path to the Internet from the repository server.) If you check Use the call home setup , this worksheet is complete.	<input type="checkbox"/> Use the call home setup <input type="checkbox"/> Use FTP	<input type="checkbox"/> Use the call home setup <input type="checkbox"/> Use FTP

Table 91. Outbound (call home) worksheet (continued)

Item or Setting	Instructions	MC1	MC2 (if applicable)
Do you use an FTP firewall?	If you are using FTP to deliver dump and trace call-home transmissions, check Yes if your dump and trace files cross a firewall that filters FTP traffic. Check No if no FTP firewall filters the transmissions. If you check Yes , complete the remaining worksheet items. If you check No , this worksheet is complete.	[] Yes [] No	[] Yes [] No
<p>FTP firewall settings: Table 92 lists the types of FTP firewall that are supported by the management console. Review Table 92 to determine the type of FTP to use if you plan to send dumps and traces to the IBM service center and if the dumps and traces pass through an FTP firewall.</p> <p>Notes:</p> <ol style="list-style-type: none"> 1. Bold options in the MC1 and MC2 columns indicate default settings. 2. Management console is abbreviated as MC in the worksheet. 			

Types of FTP proxy servers

As an alternative to an Internet SSL or VPN connection through the Internet, the management console can be set up to use FTP for sending error data to IBM. Table 92 lists the supported types of FTP proxy servers.

Table 92. Types of FTP proxy servers

Type	MC1	MC2 (if applicable)
1	Connect to the FTP proxy server and log in with USER user real host name.	Connect to the FTP proxy server and log in with USER user real host name.
2	Connect to the FTP proxy server and log in with the following settings: <ul style="list-style-type: none"> • USER fwuser • PASS fwpassword • USER user real host name 	Connect to the FTP proxy server and log in with the following settings: <ul style="list-style-type: none"> • USER fwuser • PASS fwpassword • USER user real host name
3	Connect to and log in to the FTP proxy servers. Provide the following settings: SITE real.host.name, followed by the regular USER and PASS addresses.	Connect to and log in to the FTP proxy server. Provide the following settings: SITE real.host.name, followed by the regular USER and PASS addresses.
4	Connect to and log in to the FTP proxy server. Provide the following settings: OPEN real.host.name, followed by the regular USER and PASS addresses.	Connect to and log in to the FTP proxy server. Provide the following settings: OPEN real.host.name, followed by the regular USER and PASS addresses.
5	Connect to the FTP proxy server and log in with the following settings: <ul style="list-style-type: none"> • USER user wuser@real.host.name • PASS pass fwpass 	Connect to the FTP proxy server and log in with the following settings: <ul style="list-style-type: none"> • USER user wuser@real.host.name • PASS pass fwpass

Table 92. Types of FTP proxy servers (continued)

Type	MC1	MC2 (if applicable)
6	Connect to the FTP proxy server and log in with the following settings: <ul style="list-style-type: none"> • USER fwuser real host name • PASS fwpass • USER user • PASS pass 	Connect to the FTP proxy server and log in with the following settings: <ul style="list-style-type: none"> • USER fwuser real host name • PASS fwpass • USER user • PASS pass
7	Connect to the FTP proxy server and log in with the following settings: <ul style="list-style-type: none"> • USER user real host name fwuser • PASS pass • ACCT fwpass 	Connect to the FTP proxy server and log in with the following settings: <ul style="list-style-type: none"> • USER user real host name fwuser • PASS pass • ACCT fwpass

Inbound (remote services) worksheets

Complete the inbound (remote services) worksheets to specify whether you want inbound remote services through Assist On-site (AOS), virtual private network (VPN), or the management console (MC) modem. Note that VPN and modem might be discontinued at a future version of this product.

If you do not select inbound unattended service sessions during the initial configuration, you can schedule specific unattended remote service sessions with an IBM service representative later. You can then either configure the management console to temporarily allow unattended remote services (specifying a start and stop date) or appoint a user onsite to manually enable the inbound call from the IBM service representative. This user must remain at the management console during the service session.

There are four options for unattended inbound remote service:

- To select external AOS for attended or unattended inbound remote service, see “External AOS.”
- To select embedded AOS for unattended inbound remote service, see “Embedded AOS” on page 230.
- To select virtual private network (VPN) access for unattended inbound remote service, see “Inbound VPN” on page 233.
- To select modem access for unattended inbound remote service, see “Inbound modem” on page 233. Enabling access through a modem is also recommended as a backup path for AOS or VPN access.

External AOS

External AOS is an Internet/SSL based form of remote service. It is IBM's preferred method for remote service. It gives you more security and control over communications and allows both attended and unattended remote service.

For AOS to work externally, you must provide a server to install the service configuration software. This server must be connected to your network and must allow traffic to the product and to the Internet through firewalls. IBM secure remote connections can be established with the storage system through port forwarding and SSL from this server. You can configure multiple storage systems under the same external AOS server.

Note: Storage systems at R4 or later can be remotely serviced with external AOS.

For more information about an external AOS and the correct configuration, see the IBM Redpaper™ Introduction to Assist On-Site for Storage (www.redbooks.ibm.com/redpapers/abstracts/redp4889.html)

Embedded AOS

Embedded AOS is similar to external AOS as it provides a method for secure and fast remote service. However, it does not require an external server. The management console (MC) hosts the AOS server. Note that with Embedded AOS you must configure your network such that each MC has Internet connectivity, whereas external AOS allows you to manage many MCs through a central server and interface.

Embedded AOS is a fast, broadband form of remote access. Embedded AOS is installed on the management console (MC). Select the embedded AOS option if you are unable to provide a server for external AOS.

Note: AOS embedded on the MC is available on DS8870 version 7.1 (licensed machine code level 87.0.102.0) and later. It is also available for DS8700 and DS8800 version 6.3 (licensed machine code level 86.31.142.0 or 76.31.121.0) and later. It is not available for DS8100 or DS8300.

For embedded AOS version 4 on DS8870 version 7.5 and later, allow encrypted SSL to the following servers.

Table 93. Customer firewall ports to enable embedded AOS version 4 remote services on DS8870 version 7.5 and later

GEO	Host name	IP address	Ports
EMEA	aos.uk.ihost.com	195.171.173.165	443

For embedded AOS version 3.3 on DS8870 version 7.1 (licensed machine code level 87.0.102.0) through version 7.4, and DS8700 and DS8800 version 6.3 (licensed machine code level 86.31.142.0 or 76.31.121.0) and later, allow encrypted SSL to the following servers.

Table 94. Customer firewall ports to enable embedded AOS version 3.3 remote services

GEO	Host name	IP address	Ports
Americas	aos.us.ihost.com	72.15.208.234	8200 or 80
Americas	aosback.us.ihost.com	72.15.223.61	8200 or 80 or 443
Americas	aosrelay1.us.ihost.com	72.15.223.60	8200 or 80 or 443

Notes: Management console is abbreviated as MC in the worksheet.

Table 95. Embedded AOS worksheet

Item or Setting	Instructions	MC1	MC2 (if applicable)
Allow unattended remote service sessions with embedded AOS?	Check Yes if you want to allow authorized IBM service representatives to initiate unattended remote service sessions on your storage system with embedded AOS. Check No if you do not want to allow unattended remote service sessions with embedded AOS. If you check No , you completed this worksheet.	<input type="checkbox"/> Yes (default) <input type="checkbox"/> No	<input type="checkbox"/> Yes (default) <input type="checkbox"/> No
Customer Name	Enter your customer name followed by the storage facility serial number. For example, Company Name - 75xxxx		
Access Control List	Select the appropriate group for DS8000 support. ¹		
Enable Collaboration	Ensure that this option is checked to allow remote support to use AOS.	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
Enable Session Recording	Ensure that this option is checked to enable recording a session in a log file.	<input type="checkbox"/> Enabled	<input type="checkbox"/> Enabled
Proxy information	If you checked Yes for embedded AOS, and you have a proxy, specify the proxy information.	Proxy URL: * Port: User ID: Password:	Proxy URL: ** Port: User ID: Password:

Table 95. Embedded AOS worksheet (continued)

Item or Setting	Instructions	MC1	MC2 (if applicable)
*Proxy URL for MC1:			
**Proxy URL for MC2:			
Notes:			
1. Access Control List:			
If you are in one of the following countries:			
• All EMEA countries: select IBM/BLUE/GTS/EMEA/STORAGE/DS8000 plus the following EMEA country-specific ACL.			
• Austria, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_DACH			
• Belgium, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_BENELUX			
• Denmark, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_Nordics			
• Finland, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_Nordics			
• France, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_FRANCE			
• Germany, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_DACH			
• Ireland, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_UKI			
• Italy, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_ITALY			
• Luxembourg, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_BENELUX			
• Netherlands, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_BENELUX			
• Norway, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_Nordics			
• Spain, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_SPAIN			
• Sweden, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_Nordics			
• Switzerland, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_DACH			
• UK, select only IBM/BLUE/GTS/EMEA/STORAGE/FO_UKI			
• Any other EMEA country, select only IBM/BLUE/GTS/EMEA/STORAGE/VFE_CEEMEA			
• Any other non-EMEA country, select the appropriate items:			
IBM/BLUE/STG/Americas/United States/Storage/DS8000/PFE			
IBM/BLUE/GTS/Americas/United States/Storage/DS8000/RSC			
IBM/BLUE/GTS/Americas/United States/Storage/DS8000/CE			
IBM/BLUE/GTS/Americas/Canada/Storage/DS8000/RSC			
IBM/BLUE/GTS/Asia Pacific/China/Storage/DS8000/PFE			
IBM/BLUE/GTS/Asia Pacific/China/Storage/DS8000/RSC			
IBM/BLUE/GTS/Asia Pacific/China/Storage/DS8000/CE			
IBM/BLUE/GTS/Asia Pacific/Japan/Storage/DS8000/PFE			
IBM/BLUE/GTS/Asia Pacific/Japan/Storage/DS8000/RSC			
IBM/BLUE/GTS/Asia Pacific/Japan/Storage/DS8000/CE			
IBM/BLUE/GTS/Americas/Latin Americas/Brazil/Storage/DS8000/RSC			
IBM/BLUE/GTS/Americas/Latin Americas/Argentina/Storage/DS8000/RSC			
IBM/BLUE/GTS/Americas/Latin Americas/Mexico/Storage/DS8000/RSC			
IBM/BLUE/GTS/Americas/Latin Americas/Chile/Storage/DS8000/RSC			
IBM/BLUE/GTS/Americas/Latin Americas/Colombia/Storage/DS8000/RSC			

Allowing AOS information to be displayed in PMRs in the storage system

The IBM service representative must configure more AOS details in the MC modem configuration for remote access so they display in the IBM Problem Management Record (PMR) system. This enables IBM support to recognize that a particular storage system has AOS connectivity. Use the following steps, which can be implemented through the DS Service GUI or by an IBM service representative, to display AOS information in PMRs.

1. Log in to the DS Service GUI or use either the Customer Engineer (CE) or the CUSTOMER user ID.
2. Select **Service Management > Manage Inbound Connectivity**.
3. Check **Allow unattended sessions**.

4. For embedded AOS, in the 'Modem Phone Number' field enter "AOS" and the country information from the ACL you selected in the previous section. For example, "AOS USA". For external AOS, enter "AOS" and the IP address or hostname of the management console HMC. For example "AOS 9.11.22.33".
5. Clear **Allow unattended sessions**. This option controls the answering mode of the attached modem. It does not have impact (allow or disallow) access through the network.

Note: If the customer also uses a modem for dial-in then, do not clear this option and add the modem phone number in addition to the AOS information.

6. Click **OK** to close.
7. Generate a test Call Home PMR.

Inbound VPN

If you want to enable -attended inbound remote services through a virtual private network (VPN), refer to the following information.

IBM can provide attended inbound remote support through VPN on the management console (MC). VPN provides connectivity if you do not have an inbound AOS or modem connection.

To enable VPN access for unattended inbound remote support, use the “Outbound (call home and dump/trace offload) worksheet” on page 224. Enable call home and select **By Internet VPN**. This option enables both outbound and inbound VPN access for remote services. Note that IBM might discontinue VPN connectivity in future versions of this product.

Inbound modem

If you want to enable unattended remote services through the management console (MC) modem, refer to the following information and complete the inbound modem worksheet. Note that IBM might discontinue modem connectivity in future versions of this product.

For unattended remote service through the modem, you must first use the “Outbound (call home and dump/trace offload) worksheet” on page 224 to configure the management console (MC) for call home and select the **By modem** connectivity mode.

Notes: Management console is abbreviated as MC in the worksheet.

Table 96. Inbound modem worksheet

Item or Setting	Instructions	MC1	MC2 (if applicable)
Allow unattended remote service sessions by modem?	Check Yes if you want to allow authorized IBM service representatives to initiate unattended remote service sessions on your storage system through the modem. Check No if you do not want to allow unattended remote service sessions through the modem. If you check No , this worksheet is complete.	<input type="checkbox"/> Yes (default) <input type="checkbox"/> No	<input type="checkbox"/> Yes (default) <input type="checkbox"/> No

Table 96. Inbound modem worksheet (continued)

Item or Setting	Instructions	MC1	MC2 (if applicable)
Unattended remote service session settings for modem: Complete the following section if you selected Yes to allow unattended remote service sessions by modem.			
Remote service sessions mode for modem	Check the mode that indicates when to allow unattended service sessions. Select Continuous to enable inbound remote service at any time. Select Automatic to allow inbound calls for a specified number of days that follow a failure on the storage system.	<input type="checkbox"/> Continuous (default) <input type="checkbox"/> Automatic <input type="checkbox"/> Temporary	<input type="checkbox"/> Continuous (default) <input type="checkbox"/> Automatic <input type="checkbox"/> Temporary
Number of days for Automatic mode	If you selected the Automatic mode, specify the number of days to allow an unattended service session after any failure on the storage system.		
Interval for Temporary Mode	If you selected the Temporary mode, specify the starting and ending dates of the time period when unattended service sessions are allowed. Note: This option allows control of when IBM can perform unattended service sessions. You can change the interval right before a service action takes place.		

Notification worksheets

The notification worksheets specify your preferred method of being notified about serviceable events.

Note: The IBM service representative sets up the notification process.

Use the notification worksheets to specify the settings to use when you want the storage system to notify you or other people in your organization when you have serviceable events.

There are two notification worksheets:

- SNMP trap notification worksheet
- Email notification worksheet

SNMP trap notification worksheet

Complete the SNMP trap notification worksheet to specify the setting for SNMP trap notifications.

Use the SNMP trap notification worksheet to indicate whether you want to receive Simple Network Management Protocol (SNMP) trap notifications when a management console encounters serviceable events.

Note: Remote copy status reporting for Copy Services requires SNMP for open-systems hosts.

Worksheet purpose

IBM service representatives use the information on the SNMP trap notification worksheet to customize your storage system for SNMP trap notifications.

Worksheet and instructions

You must complete Table 97 for all installations that include a management console.

Notes:

1. Bolded options in the MC1 and MC2 columns indicate default settings.
2. Management console is abbreviated as MC for the following table.

Table 97. SNMP trap notification worksheet

Item or Setting	Instructions	MC1	MC2 (if applicable)
Enable SNMP trap notifications?	Check Yes to allow the storage system to generate and send SNMP trap notifications when the system encounters problems. Check No if you do not want the storage system to send SNMP trap notifications. If you check No , this worksheet is complete.	<input type="checkbox"/> Yes <input type="checkbox"/> No	<input type="checkbox"/> Yes <input type="checkbox"/> No
SNMP trap notification settings: Complete the following section if you checked Yes to enable SNMP trap notifications. Do not use the IP address that is shown in the example in this worksheet. The IP address is only an example and does not function. Your IBM service representative can provide the correct IP address.			
SNMP trap destinations	Provide the dotted decimal addresses and community name of the destinations that are to receive SNMP traps (for example, 9.127.152.254 default). Note: If you plan to use advanced functions SNMP messaging, you must set those functions by using DS CLI.		

Email notification worksheet

Complete the email notification worksheet to specify the setting for email notifications.

Use the email notification worksheet to specify whether you want to receive email notifications when a management console encounters serviceable events.

Restriction: To receive email notifications, the management console must be connected to your LAN.

Worksheet purpose

IBM service representatives use the information on this worksheet to customize your storage system for email notifications. If you use email notifications, the notification settings are customized so that the specified people in your organization receive emails when there is general or error information to send about the storage system.

Worksheet and instructions

You must complete Table 98 for all installations that include a management console.

Notes:

1. Bold options in the MC1 and MC2 columns indicate default settings.
2. Management console is abbreviated as MC in the following table.

Table 98. Email notification worksheet

Item or setting	Instructions	MC1	MC2 (if applicable)
Enable email notifications?	Check Yes to allow the MC to generate and send emails when the system encounters problems. Check No if you do not want the MC to send email notifications. If you check No , this worksheet is completed.	<input type="checkbox"/> Yes <input type="checkbox"/> No	<input type="checkbox"/> Yes <input type="checkbox"/> No
Email notification settings: Complete the following section if you previously checked Yes (to enable email notifications).			
Host name or network address of smart relay host (Optional)	To use a smart relay host, provide the host name or network address for the smart relay host. Tip: You can enable a smart relay host if either of these conditions applies: <ul style="list-style-type: none"> • Your email is sent from a UNIX system on which you specified a mail relay or mail gateway, or • You installed a message-transfer agent on your mail server. 		

Table 98. Email notification worksheet (continued)

Item or setting	Instructions	MC1	MC2 (if applicable)
Email destinations	Provide the full email addresses where you want to receive the notifications (for example, maria@host.com). Check the notification setting that indicates the type of notifications to send to the email address. This worksheet provides spaces for three email addresses, but you can specify more, if necessary.	1. Email address: _____ Notifications: <input type="checkbox"/> Only call home problem events <input type="checkbox"/> All problem events 2. Email address: _____ Notifications: <input type="checkbox"/> Only call home problem events <input type="checkbox"/> All problem events 3. Email address: _____ Notifications: <input type="checkbox"/> Only call home problem events <input type="checkbox"/> All problem events	1. Email address: _____ Notifications: <input type="checkbox"/> Only call home problem events <input type="checkbox"/> All problem events 2. Email address: _____ Notifications: <input type="checkbox"/> Only call home problem events <input type="checkbox"/> All problem events 3. Email address: _____ Notifications: <input type="checkbox"/> Only call home problem events <input type="checkbox"/> All problem events

Power control worksheet

Complete the power control worksheet to specify the power mode for your storage system.

You can:

- Use attached IBM z Systems or S/390 hosts or S/390 hosts to power on and power off the storage system. (This option is available only if you have the remote zSeries power control feature installed.)
- Automatically power on and power off the storage system.
- Use a specified schedule to power on and power off the storage system.
- Manually power on and power off the storage system. Use the Power on/off page in the DS Storage Manager.

Worksheet purpose

IBM service representatives use the information on the power control worksheet to customize the power mode for your storage system.

Worksheet and instructions

You must complete Table 99 for all installations.

Table 99. Power control worksheet

Item or Setting	Instructions	Your information
Enable remote zSeries power control?	If you plan to use the remote zSeries power control feature, check Yes . If you check Yes , choosing zSeries power mode enables up to four z Systems or S/390 hosts to control the power-on and power-off sequences. If you check Yes , this worksheet is complete. Check No if you choose not to use the remote zSeries power control feature. If you check No , you must complete the rest of this worksheet.	<input type="checkbox"/> Yes <input type="checkbox"/> No
Disabled remote z Systems power control: Complete the following section if you checked No on whether to use the remote IBM z Systems power control.		
Power mode	Check Automatic if you want the storage system to power on automatically whenever external power is restored, if the storage system was originally on. (The Automatic power mode automatically powers on the storage system when, for example, power is restored after a power outage.) Check Scheduled if you want the storage system to power on and off according to a specified scheduled. Select Scheduled automatic to schedule the power-on and power-off actions for your storage system and enable the storage system to automatically power on if power is restored while the storage system is scheduled to be on. Check Manual if you want to manually power on and power off your storage system. You can use the Power on/off page in the DS8000 Storage Management GUI.	<input type="checkbox"/> Automatic <input type="checkbox"/> Scheduled (not automatic) <input type="checkbox"/> Scheduled automatic <input type="checkbox"/> Manual
Schedule	If you selected one of the scheduled power modes, Scheduled or Scheduled automatic , specify the power-on and power-off schedule.	

Table 99. Power control worksheet (continued)

Item or Setting	Instructions	Your information
Schedule	Check whether you prefer the storage system to have a power-on and power-off schedule that is the same every day or prefer a schedule that varies every day. Specify the on and off times for the storage system in the appropriate section.	_ Same schedule all days: On _____ Off _____ _ Varying schedule: Monday: On _____ Off _____ Tuesday: On _____ Off _____ Wednesday: On _____ Off _____ Thursday: On _____ Off _____ Friday: On _____ Off _____ Saturday: On _____ Off _____ Sunday: On _____ Off _____

Control switch settings worksheet

Complete the control switch settings worksheet to indicate whether to enable or disable a particular switch setting. Your storage complex can include one or more storage systems; so you might need to indicate your individual choice for each.

The switch settings can be enabled or disabled by indicating your choices on the worksheet that follows the setting descriptions. IBM service representative use the choices that you specify on the worksheet to set the control switches for your storage system.

16 Gb/s Fibre Channel forward error correction

This control switch provides 16 Gb Fibre Channel forward error correction (FEC). FEC allows the target to correct errors without needing a reverse channel to request retransmission of data. FEC helps prevent I/O errors from occurring and smooths the adoption of faster link speeds. Greater loss margins afforded by the use of FEC reduce the occurrence of link errors and associated service cost.

The IBM service representative can set this control switch; it cannot be set through the DS CLI or GUI.

Control-unit initiated reconfiguration settings

The control-unit initiated reconfiguration (CUIR) setting indicates whether to enable or disable subsystems. CUIR prevents loss of access to volumes in IBM z Systems environments due to incorrect path handling. This function automates

channel path management in IBM z Systems environments, in support of selected storage system service actions. CUIR relies on a combination of host software and storage system firmware. The host systems are affected during host adapter repair or I/O enclosure repair. Use the CUIR setting on the worksheet to indicate whether this option can be enabled. The CUIR setting applies to IBM z Systems and S/390 environments only.

Control unit threshold

This control unit threshold switch provides the threshold level for presenting a SIM to the operator console for controller-related errors. SIMs are always sent to the attached IBM z Systems hosts for logging to the Error Recording Data Set (ERDS). SIMs can be selectively reported to the IBM z Systems host operator console, as determined by SIM type and SIM severity. This setting applies to IBM z Systems and S/390 environments only.

Table 100. SIM presentation to operator console by severity

Selection	Severity of SIM presented to operator console
Service	Service, Moderate, Serious, and Acute (all)
Moderate	Moderate, Serious, and Acute
Serious	Serious and Acute
Acute	Acute
None	None

Table 101. SIM severity definitions

Severity	Definition
Service	No system or application performance degradation is expected in any environment.
Moderate	Performance degradation is possible in a heavily loaded environment.
Serious	A primary subsystem resource is disabled.
Acute	A major subsystem resource is disabled. Performance might be severely degraded. System or application outages might have occurred.

Device threshold

This control switch provides the threshold level for presenting a SIM to the operator console for device-related errors. Device threshold levels are the same type and severity as control unit threshold settings. Device threshold applies to IBM z Systems and S/390 environments only.

Full page protection

This control switch provides the ability to ensure that the atomicity of a database page-write is maintained.

IBM i LUN serial suffix number

Use the IBM i LUN serial suffix number switch setting only when you attach two or more storage systems that have worldwide node names (WWNNs) with the same last three digits to an AS/400 or IBM i host.

For example, the WWNN for the first storage system is 500507630BFFF958 and the WWNN for an additional storage system is 500507630AFF958. Both storage systems would present the same LUN serial number for the same LUN ID. Because the original LUN serial number is used for identification, the AS/400 does not use the LUN from the additional storage system. Specifying a unique serial number base for the additional storage system prevents this problem. The IBM service representative enters the control-switch setting for the new serial number base, which you specify for this field.

The WWNN can be found on the WWID label on the inside front left wall of the base frame, which is located near the LED indicators of the upper-primary power supply. For example, WWID: 500507630AFF99F.

Notes:

- The probability of receiving two storage systems with the same last three WWNN digits is unlikely, but possible.
- After you update the switch settings, a quiesce and resume of the storage system is required for the changes to take effect.

The IBM i LUN serial suffix number applies to IBM i and AS/400 environments only.

Lights-on fast load for the 8Gb/s host adapters

This control switch enables lights-on fast load for the 8 Gb host adapters. Enabling the control switch provides seamless microcode update to storage arrays and prevents potential loss of paths and access caused by a profusion of register state-change messages if light is dropped during fast load.

The IBM service representative sets this control switch; it cannot be set through the DS CLI or GUI.

Lights-on fast load for the 16Gb/s host adapters

This control switch enables lights-on fast load for the 16 Gb host adapters. Enabling the control switch provides seamless microcode update to storage arrays and prevents potential loss of paths and access caused by a profusion of register state-change messages if light is dropped during fast load.

The IBM service representative sets this control switch; it cannot be set through the DS CLI or GUI.

Media threshold

This control switch provides the threshold level for presenting a SIM to the operator console for media-related errors. Media threshold levels are the same type and severity as control unit threshold settings. Media threshold applies to IBM z Systems and S/390 environments only.

Present SIM data to all hosts

Service Information Messages (SIMs) are offloaded to the first I/O request. The SIMs are directed to each logical subsystem in the storage system if the request is device or control unit related. The SIMs are offloaded to the individual logical volume when the request is media-related. This control switch determines whether

SIMs are sent to all, or to only the first attached IBM z Systems LPAR and makes an I/O request to the logical system or logical volume. This setting applies to IBM z Systems and S/390 environments only.

IBM z Systems high-performance FICON enhanced buffer management

This control switch provides IBM z Systems high-performance FICON enhanced buffer management. It allows the channel to send larger (>64 Kb) writes to the storage system without requiring the use of the Transfer_Ready command which, at long distances, increases the latency in the I/O operation.

The IBM service representative can set this control switch; it cannot be set through the DS CLI or the GUI.

Use Table 102 to enter the appropriate response into the information column.

Table 102. Control switch settings worksheet

Control Switch Setting	Default	Your information
16 Gb/s Fibre Channel forward error correction	1 (Enable)	<input type="checkbox"/> true = Enable <input type="checkbox"/> false = Disable
Control unit threshold	2	Present SIM to operator console for the following severities: <input type="checkbox"/> 0 = Service, Moderate, Serious, and Acute (all) <input type="checkbox"/> 1 = Moderate, Serious, and Acute <input type="checkbox"/> 2 = Serious and Acute <input type="checkbox"/> 3 = Acute
CUIR support	0 (Disable)	<input type="checkbox"/> true = Enable CUIR support <input type="checkbox"/> false = Disable CUIR support
Device threshold	2	Present SIM to operator console for the following severities: <input type="checkbox"/> 0 = Service, Moderate, Serious, and Acute (all) <input type="checkbox"/> 1 = Moderate, Serious, and Acute <input type="checkbox"/> 2 = Serious and Acute <input type="checkbox"/> 3 = Acute
Full page protection	1 (Enable)	<input type="checkbox"/> true = Enable <input type="checkbox"/> false = Disable
IBM i LUN Serial Suffix number - AS/400 LUN Serial Suffix number	0 (Off)	_____ Enter the last three digits of the storage system worldwide node name (WWNN). ¹ <input type="checkbox"/> 0 = Off (use last three digits of WWNN) <input type="checkbox"/> _____ (enter three numeric digits to create a unique identifier)
Lights-on fast load for the 8Gb/s host adapters	0 (Disable)	<input type="checkbox"/> true = Enable <input type="checkbox"/> false = Disable
Lights-on fast load for the 16Gb/s host adapters	1 (Enable)	<input type="checkbox"/> true = Enable <input type="checkbox"/> false = Disable
Media threshold	2	Present SIM to operator console for the following severities: <input type="checkbox"/> 0 = Service, Moderate, Serious, and Acute (all) <input type="checkbox"/> 1 = Moderate, Serious, and Acute <input type="checkbox"/> 2 = Serious and Acute <input type="checkbox"/> 3 = Acute
Present SIM data to all hosts	0 (Disable)	<input type="checkbox"/> true = Enable (all hosts) <input type="checkbox"/> false = Disable (host issuing start I/O)

Table 102. Control switch settings worksheet (continued)

Control Switch Setting	Default	Your information
IBM z Systems high-performance FICON enhanced buffer management	1 (Enable)	[] true = Enable [] false = Disable
<p>Note: The WWNN can be determined only by reading the WWNN label on the front left inside wall of the base frame after it is unpacked. Only then can you determine whether there is a duplication of the last three digits with an existing storage system. If you cannot wait for that to occur, enter three numeric digits to create a unique identifier.</p>		

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.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
3-2-12, Roppongi, Minato-ku, Tokyo
106-8711 Japan*

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.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Copyright and trademark information website (www.ibm.com/legal/copytrade.shtml).

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

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

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Index

Numerics

- 2-bay racks 39
- 2.5-inch 38
- 2244 Model PAV 127
- 2244 Model PTC 129
- 239x Model LFA 127
- 3.5-inch 38
- 4-bay racks 39
- 4-port HA 39
- 8-port HA 39

A

- accessibility
 - features 207
- acclimation 139
- activating licenses 202
- adapters 182
 - configuration rules 109
 - Fibre Channel host 105, 106
- advisor tool 68, 69
- air circulation 173, 176
 - intake and exhaust 173
- air quality 174
- algorithms 45
- allocation methods 24
- AS/400 LUN
 - control switch settings 187
- Attachment to IBM z Systems 187
- auto-performance rebalance 53
- auto-rebalance 51
- auto-rebalancing 53
- automatic 53
- automatic data migration with Easy Tier 53
- auxiliary volumes 51
- availability features 38

B

- balancing the I/O load 42
- battery assemblies 114
- battery service modules
 - feature codes 115
- BSMI certificate 116
- BTU 171

C

- cable
 - configuration rules 109
 - cutout locations 146
 - disk drive 97
 - drive cables, feature codes 97
 - feature codes, Fibre Channel
 - cable 106
 - Fibre Channel host adapter 105
 - I/O adapters 104
 - I/O cables 105

- cable (*continued*)
 - installation 146
 - overhead cable 146
 - RIO-G 105
 - top exit bracket 146
- cables
 - drive 97
- cache 111, 112
- canceling migration 67
- capacity
 - calculating physical and effective 99
 - exhaust 171
 - floor load rating 150
- caution notices vii, viii
- CCW, channel command words 43
- certificate, BSMI 116
- circuit breakers
 - high-voltage 167
 - low-voltage 167
- CKD, count key data storage 43
- clearances required for service 152
- CLI, command-line interface 26
- clusters, RAID disk groups 42
- CoD 38
- cold demote 53
- comments, sending xxi
- communication requirements, host
 - attachment 182
- company information 185
 - worksheet 217
- configuration
 - battery service modules 115
 - disk drive cables 97
 - DS Storage Manager 26
 - I/O (RIO-G) cables 105
 - processor memory 111
 - reconfiguration 26
- configuration control indicators 87
- configuration overview 18
- Configuration overview 20
- configuration rules
 - device adapters 107
 - flash interface cards 107
 - host adapters and cables 109
 - I/O adapter 107
 - I/O enclosures 107
 - management consoles 90
 - Standby CoD disk drives 97
 - storage devices 97
 - storage enclosures 97
 - system memory 112
- connectivity
 - I/O enclosures 104
- consolidating storage 42
- containers, shipping 139
- contamination information 175
- control switch settings 187
 - CUIR 239
 - serial suffix number 239
 - SIMs 239
 - threshold settings 239

- control switch settings (*continued*)
 - worksheet 239
- control unit threshold 187
- conventions
 - terminology xvii
 - typefaces xvii
- cooling 173, 176
- Copy Services
 - considerations 72
 - disaster recovery 81
 - licensed functions 84
 - overview 72
 - point-in-time function 129
 - remote
 - overview 130
 - z/OS Global Mirror 132
 - SNMP trap notification work sheet 234
- corrosive gasses and particulates 174
- count key data storage 43
- CUIR, control-unit initiated
 - reconfiguration 187

D

- danger notices vii, xiii
- data
 - securing 86
- data migration
 - selecting method 192
- data movement daily report 68
- data placement 42
- DC-UPS 114
- description
 - VMware 69
- description of Easy Tier 51
- description of EAV 46
- device adapters 96
 - configuration rules 107
 - feature codes 96
- device driver, subsystem 41
- Device threshold 187
- dimensions
 - storage system, installed 151
- disaster recovery
 - Copy Services 81
- disk drive
 - cable 90
 - cables 97
- disk drive module
 - maintenance policy 38
- disk drive sets 90
- disk drives 38
 - subsystem device driver 41
- disk enclosures 90
 - fillers 90
- disk intermix
 - configuration indicators 87
- Disk Manager
 - monitoring performance 28
- DNS settings 219

- drive enclosures
 - See* standard drive enclosures
- drive set
 - capacity 99
- drive sets 90, 92
- drives
 - cable 97
 - capacity calculation 99
- DS command-line interface 26
- DS Storage Manager 26
- DS8000
 - architecture 3
 - implementation 3
- DS8800 213
- DS8800 to DS8870
 - model conversion 213
- DS8870 213
- DSFA, Disk Storage Feature
 - Activation 202
- dynamic expansion
 - volume 50
- dynamic volume expansion 50

E

- EAM 44
- earthquake preparedness 179
- earthquake resistance kit
 - required preparation for 153, 154, 159
- Easy Tier 24, 53
 - application controls 59
 - Application for z Systems 58
 - feature codes 128
 - heat map transfer 64
 - licensed functions 128
 - manual mode 61
 - overview 51
 - pool merge 61
 - Storage Tier Advisor Tool
 - data movement daily report 68
 - workload categorization 68
 - workload skew 68
 - volume migration 61
- Easy Tier application controls 59
- Easy Tier Application for z Systems 58
- Easy Tier Server
 - feature codes 128
- EAV CKD
 - 1 TB z Systems CKD 46
 - 3390 Model A 46
 - cylinder size 46
- email notification worksheet 235
- enclosure fillers
 - feature codes 96
- encryption
 - overview 86
 - planning for 195, 196
- environment 171
 - air circulation 173
 - operating requirements 173
- ePLD
 - See* extended power-line disturbance
- EPLD
 - See* extended power line disturbance
- ESE and auxiliary volumes 61
- ESE volumes 51

- Ethernet settings 219
- exhaust 171
- expansion model position
 - configuration indicators 87
- extended address volumes
 - overview 46
- extended power line disturbance 115
- extended power-line disturbance 116
- Extended Remote Copy (XRC) (*see* z/OS Global Mirror) 132
- external power cords 90

F

- failover and fallback 81
- FB, fixed block 43
- FDE 38
- Feature Activation (DSFA), IBM Disk Storage 202
- feature codes
 - See also* features
 - additional setup options 115
 - battery service modules 115
 - device adapters 96
 - drive cables 97
 - drive sets 92
 - Easy Tier 128
 - Easy Tier Server 128
 - enclosure fillers 96
 - extended power line disturbance 115
 - extended power-line disturbance 116
 - features, other configuration 115
 - Fibre Channel cable 106
 - Fibre Channel host adapters 105
 - flash cards 92
 - flash enclosures 95
 - FlashCopy licensed function 129
 - FlashCopy SE 130
 - hardware management console,
 - external power cords 90
 - I/O (RIO-G) cables 105
 - I/O adapter 104
 - I/O enclosures 104
 - I/O Priority Manager 132
 - IBM HyperPAV 127
 - management console 89
 - memory 111
 - optional 115
 - ordering optional features 88
 - overhead cable management 145
 - parallel access volume 127
 - physical configuration 88
 - power control 116
 - power cords 112, 113
 - power features 115
 - processors 111
 - remote mirror and copy 131
 - remote zSeries power control 116
 - setup 115
 - shipping 115
 - shipping and setup 115
 - shipping weight reduction 117
 - standard drive enclosures 95
 - thin provisioning 136, 137
 - z/OS Metro/Global Mirror
 - Incremental Resync 133
 - zSeries 116

- features
 - input voltage
 - about 114
- feedback, sending xxi
- Fibre Channel
 - host adapters 105
 - host adapters feature codes 105
 - host attachment 182
 - open-systems hosts 40
 - SAN connections 40
- Fibre Channel cable
 - feature codes 106
- fire suppression 179
- fixed block storage 43
- flash cards 38, 92
 - overview 91
- flash copy 82
- flash drives 38
- flash drives, feature codes 92
- flash enclosure
 - overview 91
- flash enclosures
 - feature codes 95
- flash interface cards 96
 - configuration rules 107
- flash RAID adapters 96
- FlashCopy
 - feature codes for licensed
 - function 129
 - Multiple incremental FlashCopy 75
 - space-efficient 49
 - FlashCopy SE 49
 - feature codes 130
- floor and space requirements 141
- floor load 150
- force option 44, 50

G

- Global Mirror 81, 82

H

- HA intermix 39
- hardware features 31
- hardware management console 90
- hardware planning 31
- high-voltage
 - installations 167
- homogeneous pools 51
- homologation 249
- host adapters
 - Fibre Channel 105
- host attachment
 - overview 39
- host systems
 - communication requirements 182
- hot spot management 51
- how to order using feature codes 88
- HyperPAV 127

I

- I/O adapter
 - configuration rules 107
 - features 104

- I/O cable
 - configuration 105
- I/O drawer 39
- I/O enclosures 104
 - configuration rules 107
 - feature codes 104
- I/O load, balancing 42
- I/O plug order 39
- I/O Priority Manager
 - feature codes 132
- IBM Disk Storage Feature Activation (DSFA) 202
- IBM DS8000 Storage Tier Advisor Tool
 - overview 68
- IBM HyperPAV
 - feature code 127
- IEC 60950-1 vii
- implementation, RAID 22, 36
- inbound (remote services) work sheet 229
- initialization 47
- input voltage 114
 - configuration rules 115
- input voltage requirements 167
- installation
 - air circulation 176
 - components with shipment 211
 - external management console 179
 - nonraised floor 142
 - raised floor 142
- installation planning 139
- IOPS 51

L

- labels, safety information vii
- LFF 38
- licensed function
 - FlashCopy 129
- licensed functions 202
 - Copy Services 84
 - Easy Tier 128
 - HyperPAV 127
 - parallel access volumes 127
 - thin provisioning 135
 - z/OS Distributed Data Backup 134, 135
- licenses
 - Disk Storage Feature Activation (DSFA) 202
 - function authorization
 - documents 202
- limitations 69
- logical subsystems
 - overview 24
- logical volumes 44
- low-voltage
 - installations 167
- LSS, logical subsystem 24
- LUN
 - calculation 45
 - control switch settings 187

M

- machine types 2

- maintenance policy
 - disk drive module 38
- managed allocation 24
- management console 179
 - configuration rules 90
 - ESSNI 89
 - feature codes 89
 - HMC 89
 - HMC remote access 89
 - internal and external 89
 - multiple storage systems 36
 - network settings 185
 - network settings worksheet 219
 - overview 36
 - planning installation 179
 - power cords, external 89
 - rack specifications, external 180
 - SNMP trap notification 234
 - TCP/IP ports 181
- management consoles
 - feature codes 88
- management interfaces 26
- manual mode using Easy Tier 61
- Media threshold 187
- memory
 - feature codes 111
 - system memory 111
- memory, configuration rules 112
- migrating data
 - selecting method 192
- migration 53
 - canceling 67
 - pausing 67
 - resuming 67
- migration considerations 69
- mobile app 28
- model conversion 213
 - DS8800 to DS8870 213
- monitoring
 - with Easy Tier 64
- Multiple incremental FlashCopy 75

N

- network settings 185
- new features xxiii
- nodes 182
- noise level 171
- notices
 - caution viii
 - danger xiii
 - safety vii
- notification of service events
 - email 235
 - SNMP trap 234
- notification settings
 - methods 187

O

- obtaining activation codes 202
- operating environment
 - power on 173
 - while in storage 173
 - with power on or off 173
- outbound work sheet 224

- overview 69
 - host attachment 39

P

- parallel access volume 127
 - feature codes 127
- parallel access volume (PAV)
 - understanding static and dynamic 71
- pass-through 82
- pausing migration 67
- PAV (parallel access volumes) 71
- performance data 68
- performance gathering
 - Disk Manager 28
- physical capacity
 - remote mirror and copy 131
- physical configuration
 - drive capacity 99
 - drive enclosures 90
 - drives 90
 - extended power line disturbance 115
 - flash enclosures 95
 - I/O adapter features 104
 - I/O cable 104
 - I/O enclosures 104
 - input voltage of power supply 114
 - keyboard 90
 - management console 90
 - management consoles 88, 89
 - power cords 112, 113
 - power features 115
 - processors 111
 - standard drive enclosures 95
 - Standby CoD drives 92
- physical configuration of DS8000
 - remote zSeries power control
 - feature code 116
- planning
 - activating
 - full-disk encryption 197
 - disk encryption
 - activating 197
 - planning 197
 - earthquake resistance kit site
 - preparation 153, 154, 159
 - encryption 195, 196
 - environment requirements 173
 - external management console 179
 - floor load 150
 - full-disk encryption activation 197
 - IBM Full Disk Encryption 195, 196
 - IBM Security Key Lifecycle Manager 195
 - IBM Tivoli Key Lifecycle Manager 195, 196
 - model conversion 213
 - network and communications 181
 - operating environment, power 174
 - power connector 167
 - safety 178
 - storage complex setup 185
 - weight 150
- point-in-time copy function 129
- pool rebalancing 51
- power
 - consumption 171

- power (*continued*)
 - extended power line disturbance feature 115
 - operating environment, off 174
 - operating environment, on 173
 - outlet requirements 166
 - remote control 116
 - remote on/off 116
- power connector
 - requirements 167
 - specifications 167
- power control worksheet 237
- power cords 112
 - feature codes 113
 - hardware management console, external 90
 - management consoles 89
 - power connector 167
- power features
 - configuration rules 115
- power frequencies 167
- power supply
 - input voltage of 114
- power control settings 187
- Present SIM data to all hosts 187
- processor
 - feature codes 111
 - memory (cache) 111
- processors
 - feature codes 111
- publications
 - ordering xxi
 - product xvii
 - related xvii

Q

- quick initialization 47

R

- rack, external management console 180
- RAID
 - disk groups 42
 - implementation 22, 36
 - RAID 10 overview 23, 37
 - RAID 5 overview 23, 37
 - RAID 6 overview 23, 37
 - RAID overview 22, 36
- raised floors
 - cutting tiles for cables 146
- rank depopulation 51, 61
- redundancy
 - management consoles 89
- remote mirror and copy
 - feature codes 131
- remote mirror for z Systems (see z/OS Global Mirror) 132
- remote power control 116, 182
- remote support
 - connections 181
 - inbound 229
 - settings 186
- remote zSeries power control 116
 - See physical configuration of DS8000

- replication
 - copy services functions 28
- requirements
 - external management console installation 179
 - floor and space 141
 - floor load 150, 151
 - host attachment communication 182
 - input voltage 167
 - loading dock 140
 - modem 181
 - planning network and communications 181
 - power connectors 167
 - power outlets 166
 - receiving area 140
 - service clearance 152
 - space 151
- resource groups
 - copy services 82
- resuming migration 67
- RGL 82
- RIO-G cable 105
- rotate capacity 24
- rotate volumes 24

S

- safety 178
 - earthquake preparedness 179
 - earthquake resistance kit 153, 154, 159
 - fire suppression 179
 - information labels vii
 - notices vii
 - operating environment 179
 - power outlets 166
 - temperature and cooling 179
- SAN
 - connections with Fibre Channel adapters 40
- SAS 38
- SAS enterprise and NL SAS 51
- SATA 38
- scenarios
 - adding storage 204
- scope limiting
 - disaster recovery 82
- SDD 41
- security
 - best practices
 - service accounts 198
 - user accounts 197
- serial number setting 187
- service clearance requirements 152
- service events, outbound notification of 224
- Service Information Messages 187
- SFF 38
- shipments
 - authorized service components 212
 - container weight, dimensions 139
 - hardware, software 211
 - loading ramp 140
 - media 212
 - planning for receipt 140
 - receiving area 140

- shipments (*continued*)
 - reducing weight of 116
 - requirements
 - loading ramp 140
 - weight reduction feature code 117
- shipping containers 139
- shipping weight reduction 116
 - feature code 117
- SIM 187
- Simple Network Management Protocol (SNMP)
 - trap notification work sheet 234
- slot plug order 39
- SNMP
 - worksheet 234
- space requirements 151
- specifications
 - power connectors 167
 - specifications, non-IBM rack 180
- standard drive enclosures 90
 - feature codes 95
- standards
 - air quality 174
- Standby CoD
 - configuration indicators 87
 - disk drives 90
 - outbound work sheet 224
- Standby CoD disk drives 97
- Standby CoD drives 92
- statement of limited warranty 209
- storage area network
 - connections with Fibre Channel adapters 40
- storage features
 - configuration rules 97
 - drives and enclosures 90
- storage image
 - cooling 176
- storage system
 - service clearances 152
- Storage Tier Advisor Tool
 - Easy Tier
 - data movement daily report 68
 - workload categorization 68
 - workload skew 68
- storage-enclosure fillers 95
- storage, consolidating 42
- subsystem device driver (SDD) 41
- System i
 - control switch settings 187
- system summary report 68

T

- T10 DIF
 - ANSI support 43
 - Data Integrity Field 43
 - FB LUN 43
 - FB LUN protection 43
 - FB volume 43
 - Linux on z Systems 43
 - SCSI end-to-end 43
 - standard protection 43
 - z Systems 43
- Taiwan BSMI certificate 116
- terminology xvii
- thermal load 171

- thin provisioning 135
 - ESE capacity controls for 136
 - feature code 136
 - feature codes 137
- thin provisioning and Easy Tier 51
- three tiers 51, 53
- tiles, perforated for cooling 176
- Tivoli Storage Productivity Center 28, 36
 - copy services 28
 - Replication 28
- top exit
 - bracket 143, 146
 - measurements 143, 146
 - overhead cable management 143, 146
 - placement 143, 146
- top exit bracket
 - feature codes 145
- trademarks 246
- TSE volumes 51

U

- understanding fixed block (FB)
 - architecture 43
- understanding logical volumes 44
- user interfaces 26

V

- VMware
 - Array Integration support 69
 - restrictions 69
- volume capacity
 - overview 45
- volume deletion
 - force option 50
 - safe option 50
- volume migration 51
- volume rebalancing 51
- volumes
 - allocation 24, 44
 - deletion 44
 - force option 44
 - modification 44

W

- warm demote 53
- warranty 209
- websites xvii
- weight
 - floor load capacity 150
 - reducing shipment 116
 - feature code 117
 - storage system, installed 151
- weight and dimensions
 - shipping container 139
- work sheet
 - call home 224
 - trace dump 224
- workload categorization 68
- workload skew 68
- worksheet
 - company information 217
 - control switch settings 239
 - email notification 235

- worksheet (*continued*)
 - inbound remote support 229
 - management console network
 - settings 219
 - power control 237
 - SNMP trap notification 234
- worksheets
 - IBM-provided equipment 211
 - WWID, switch settings 187

X

- XRC (Extended Remote Copy) (see z/OS Global Mirror) 132

Z

- z Systems
 - HyperPAV 127
 - parallel access volume 127
 - remote power control settings 187
- z Systems hosts
 - FICON attachment overview 40
- z/OS Distributed Data Backup
 - licensed functions 134, 135
- z/OS Global Mirror 132
- z/OS Metro/Global Mirror Incremental Resync
 - feature codes 133



Printed in USA

GC27-4209-11

