


IBM DS8880
Version 8 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 213 and "Notices" on page 211.

This edition applies to version 8, release 5 of IBM DS8000 and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces GC27-8525-15.

© **Copyright IBM Corporation 2004, 2018.**

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

Contents

About this book	v
Who should use this book	v
Conventions and terminology	v
Publications and related information	v
IBM Publications Center	ix
Sending comments	x

Summary of changes	xi
---------------------------	-----------

Chapter 1. Overview

Machine types overview	4
Hardware	5
System types	6
Storage enclosures	21
Management console	22
Ethernet switches	22
Processor nodes	22
I/O enclosures	22
Power	23
Functional overview	23
Logical configuration	27
Logical configuration with DS8000 Storage Management GUI	27
Logical configuration with DS CLI	29
RAID implementation	31
Logical subsystems	33
Allocation methods	33
Management interfaces	35
DS8000 Storage Management GUI	35
DS command-line interface	35
DS Open Application Programming Interface	36
RESTful API	36
IBM Spectrum Control	37
IBM Copy Services Manager	37
DS8000 Storage Management GUI supported web browsers	38

Chapter 2. Hardware features

Storage complexes	46
Management console	46
Hardware specifics	46
Storage system structure	46
Disk drives and flash drives	47
Drive maintenance policy	47
Host attachment overview	48
I/O load balancing	50
Storage consolidation	50
Count key data	51
Fixed block	51
T10 DIF support	51
Logical volumes	52
Allocation, deletion, and modification of volumes	52
LUN calculation	53
Extended address volumes for CKD	54
Quick initialization	55

Chapter 3. Data management features

Transparent cloud tiering	57
Dynamic volume expansion	59
Count key data and fixed block volume deletion prevention	59
Thin provisioning	59
Extent Space Efficient (ESE) capacity controls for thin provisioning	60
IBM Easy Tier	61
VMware vStorage API for Array Integration support	64
Performance for IBM Z	66
Copy Services	67
Disaster recovery through Copy Services	76
Resource groups for Copy Services scope limiting	77
Comparison of Copy Services features	79
I/O Priority Manager	80
Securing data	81

Chapter 4. Planning the physical configuration

Configuration controls	83
Determining physical configuration features	83
Management console features	84
Primary and secondary management consoles	84
Configuration rules for management consoles	85
Storage features	85
Storage enclosures and drives	85
Storage-enclosure fillers	88
Device adapters and flash RAID adapters	89
Drive cables	89
Configuration rules for storage features	91
Physical and effective capacity	93
I/O adapter features	98
I/O enclosures	99
Fibre Channel (SCSI-FCP and FICON) host adapters and cables	99
zHyperLink adapters and cables	101
Feature codes for Transparent cloud tiering adapters	102
Feature codes for flash RAID adapters	102
Configuration rules for I/O adapter features	103
Processor complex features	108
Feature codes for processor licenses	109
Processor memory features	109
Feature codes for system memory	109
Power features	110
Power cords	110
Input voltage	112
Direct-current uninterruptible-power supply	112
Configuration rules for power features	113
Other configuration features	113
Extended power line disturbance	113
BSMI certificate (Taiwan)	113
Shipping weight reduction	114

Chapter 5. Planning use of licensed functions	115		
Licensed function indicators	115		
License scope	115		
Ordering licensed functions.	116		
Rules for ordering licensed functions	117		
Base Function license	118		
Database Protection	119		
Encryption Authorization	119		
IBM Easy Tier	120		
I/O Priority Manager	120		
Operating environment license	120		
Thin provisioning	120		
z-synergy Services license	121		
High Performance FICON for z Systems	121		
IBM HyperPAV.	122		
Parallel Access Volumes	122		
Transparent cloud tiering	122		
z/OS Distributed Data Backup	122		
zHyperLink	123		
Copy Services license.	123		
Remote mirror and copy functions	124		
FlashCopy function (point-in-time copy)	124		
Safeguarded Copy.	124		
z/OS Global Mirror	125		
z/OS Metro/Global Mirror Incremental Resync	125		
Copy Services Manager on the Hardware Management Console license	125		
Chapter 6. Delivery and installation requirements	127		
Delivery requirements	127		
Acclimation	127		
Shipment weights and dimensions	128		
Receiving delivery.	129		
Installation site requirements	130		
Planning for floor and space requirements.	131		
Planning for power requirements.	155		
Planning for environmental requirements	163		
Planning for safety	169		
Planning for network and communications requirements	169		
Chapter 7. Planning your storage complex setup	173		
Company information	173		
Management console network settings	173		
Remote support settings.	174		
Notification settings	175		
Power control settings	175		
Control switch settings	175		
Chapter 8. Planning data migration	179		
Selecting a data migration method	180		
Chapter 9. Planning for security	183		
Planning for data encryption	183		
		Planning for encryption-key servers	183
		Planning for key lifecycle managers	184
		Planning for full-disk encryption activation	185
		Planning for user accounts and passwords	185
		Managing secure user accounts	185
		Managing secure service accounts	186
		Planning for NIST SP 800-131A security conformance.	186
		Chapter 10. License activation and management.	189
		Planning your licensed functions	189
		Activation of licensed functions	190
		Activating licensed functions	190
		Scenarios for managing licensing.	191
		Adding storage to your machine	191
		Managing a licensed feature	192
		Appendix A. Accessibility features	193
		Appendix B. Warranty information	195
		Appendix C. IBM equipment and documents	197
		Installation components	197
		Customer components	198
		Service components	198
		Appendix D. DS8870 to DS8880 model conversion	199
		DS8870 to DS8880 model conversion summary	199
		Checking your preparations	201
		Removing data, configuration, and encryption	202
		Completing post-conversion tasks	202
		Appendix E. New features for models 980, 981, 982, 98B, 98E, and 98F	205
		Support for High Performance Flash Enclosures Gen2	205
		zHyperLink	206
		Transparent cloud tiering adapters	207
		Appendix F. Customization worksheets	209
		Notices	211
		Trademarks	212
		Homologation statement	213
		Safety and environmental notices.	213
		Safety notices and labels.	213
		Environmental notices	222
		Electromagnetic compatibility notices	222
		Index	227

About this book

This book describes how to plan for a new installation of DS8880. 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 (http://www.ibm.com/support/knowledgecenter/ST5GLJ_8.1.0/com.ibm.storage.ssic.help.doc/f2c_securitybp.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>DS8882F Introduction and Planning Guide</i>	This publication provides an overview of the DS8882F, the latest storage system in the DS8000 series. The DS8882F provides the new model 983. This publication provides an overview of the product and technical concepts for DS8882F.	V8.5.0 GC27-9259-00
<i>DS8880 Introduction and Planning Guide</i>	This publication provides an overview of the product and technical concepts for DS8880. It also describes the ordering features and how to plan for an installation and initial configuration of the storage system.	V8.5.0 GC27-8525-16 V8.4.0 GC27-8525-15 V8.3.3 GC27-8525-14 V8.3.1 GC27-8525-13 V8.3.0 GC27-8525-12 V8.2.3 GC27-8525-11 V8.2.1 GC27-8525-09 V8.2.0 GC27-8525-07 V8.1.1 GC27-8525-06 V8.1.0 GC27-8525-05 V8.0.1 GC27-8525-04 GC27-8525-03 V8.0.0 GC27-8525-02
<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.	V7.5.0 GC27-4209-11 V7.4.0 GC27-4209-10 V7.3.0 GC27-4209-09 V7.2.0 GC27-4209-08 V7.1.0 GC27-4209-05 V7.0.0 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.	V6.3.0 GC27-2297-09 V6.2.0 GC27-2297-07

Table 1. DS8000 series product publications (continued)

Title	Description	Order number
<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.	V8.5.0 SC27-8526-09 V8.3.3 SC27-8526-08 V8.3.1 SC27-8526-07 V8.3.0 SC27-8526-06 V8.2.3 SC27-8526-05 V8.2.2 SC27-8526-04 V8.2.0 SC27-8526-03 V8.1.1 SC27-8526-02 V8.1.0 SC27-8526-01 V8.0.0 SC27-8526-00 V7.5.0 GC27-4212-06 V7.4.0 GC27-4212-04 V7.3.0 GC27-4212-03 V7.2.0 GC27-4212-02 V7.1.0 GC27-4212-01 V7.0.0 GC27-4212-00 V6.3.0 GC53-1127-07
<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 hosts.	V8.0.0 SC27-8527-00 V7.5.0 GC27-4210-04 V7.4.0 GC27-4210-03 V7.2.0 GC27-4210-02 V7.1.0 GC27-4210-01 V7.0.0 GC27-4210-00 V6.3.0 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>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.	V7.3.0 GC27-4211-03 V7.2.0 GC27-4211-02 V7.1.0 GC27-4211-01 V7.0.0 GC35-0516-10 V6.3.0 GC35-0516-10
<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.	V1.3 SC27-9235-00 V1.2 SC27-8502-02 V1.1 SC27-8502-01 V1.0 SC27-8502-00

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

Title	Location
<i>Warranty Information for DS8000 series</i>	http://www.ibm.com/support/docview.wss?uid=ssg1S7005239
<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

Table 2. DS8000 series warranty, notices, and licensing publications (continued)

Title	Location
<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>	http://www-03.ibm.com/software/sla/sladb.nsf/pdf/ipla/\$file/ipla_en.pdf http://www-304.ibm.com/jct03001c/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 websites

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

Table 3. 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, documentation, 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 Redbooks [®] website(www.redbooks.ibm.com/)	Find technical information developed and published by IBM International Technical Support Organization (ITSO).
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.

Table 3. DS8000 series related websites (continued)

Title	Description
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.
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.
IBM Spectrum Control [™] online product documentation in IBM Knowledge Center (www.ibm.com/support/knowledgecenter)	This online documentation provides information about IBM Spectrum Control, 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.
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 displays with the DS CLI ver command when you specify the -l parameter.
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 displays with the DS CLI ver command when you specify the -l 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 displays with the DS CLI ver command when you specify the -l parameter.
DS8880 Code Bundle Information website(www.ibm.com/support/docview.wss?uid=ssg1S1005392)	Find information about code bundles for DS8880. The version of the currently active installed code bundle displays with the DS CLI ver command when you specify the -l parameter.

IBM Publications Center

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

Procedure

The IBM Publications Center website (ibm.com/shop/publications/order) offers customized search functions to help you find the publications that you need. You can view or download publications at no charge.

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 ibmkc@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

DS8000 Version 8, Release 5 introduces the following new features.

Version 8.5

This table provides the current technical changes and enhancement to the IBM DS8000 as of September 8, 2018. Changed and new information is indicated by a vertical bar (|) to the left of the change.

Note: The new system type DS8882F model 983 is covered in a separate publication. For information on the DS8882F model 983, refer to the *DS8882F Introduction and Planning Guide* (GC27-9259-00).

Function	Description
Safeguarded Copy	See “Safeguarded Copy” on page 124 for more information.
Client-side encryption for transparent cloud tiering	See “Transparent cloud tiering” on page 57 for more information.
Releasing space on CKD volumes that use thin provisioning	See “Thin provisioning” on page 59 for more information.
Easy Tier enhancements	Easy Tier infrastructure is enhanced to improve speed and efficiency. See “IBM Easy Tier” on page 61 for more information about Easy Tier.

Chapter 1. Overview

The IBM DS8880 is a high-performance, high-capacity storage system that supports continuous operation, data security, and data resiliency. For high-availability, the hardware components are redundant.

DS8880 system types with frames included add a base frame and expansion frame to the 283x machine type family, and the 533x all-flash machine type family.

Note: The modular rack-mountable system DS8882F model 983 is covered in a separate publication. For information on the DS8882F model 983, refer to the *DS888F Introduction and Planning Guide* (GC27-9259-00).

- 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 HMC is a small form factor computer and uses a keyboard and monitor that are stored in the base frame. An optional secondary HMC is also available in the base frame. A secondary HMC can provide high-availability, particularly for important processes such as encryption, Copy Services, and the HMC storage management functions.
- Depending on the system configuration, you can add up to four expansion frames to the storage system. Only the first expansion frame contains I/O enclosures, which provide more host adapters, device adapters, and High Performance Flash Enclosure Gen2 flash RAID adapters.

The DS8880 features five system types with frames: DS8884, DS8884F, DS8886, DS8886F, and DS8888F. The DS8884 (machine type 283x models 984 and 84E) is an entry-level, high-performance storage system. The DS8884F (machine type 533x model 984) is an entry-level, high-performance storage system featuring all High Performance Flash Enclosures Gen2. The DS8886 is a high-density, high-performance storage system with either single-phase power (machine type 283x models 985 and 85E) or three-phase power (machine type 283x models 986 and 86E). The DS8886F is a high-density, high-performance storage system with either single-phase power (machine type 533x models 985 and 85E) or three-phase power (machine type 533x models 986 and 86E) featuring all High Performance Flash Enclosures Gen2. The DS8888F (machine type 533x models 988 and 88E) is a high-performance, high-efficiency storage system featuring all High Performance Flash Enclosures Gen2.

Note: Previously available DS8880 models (980, 98B, 981, 98E, 982, 98F) are still supported, but not covered in this version documentation. For information on models not documented here, refer to previous version documentation (GC27-8525-06).

- The DS8884 (machine type 283x models 984 and 84E) storage system includes 6-core (12-core with zHyperLink support) processors and is scalable with up to 96 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives, up to 768 standard drives, up to 256 GB system memory, and up to 64 host adapter ports. The DS8884 includes a base frame (model 984), up to two expansion frames (model 84E), and a 40U capacity in each frame.
- The DS8884F (machine type 533x model 984) storage system includes 6-core (12-core with zHyperLink support) processors and is scalable with up to 192

Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives, up to 256 GB system memory, and up to 64 host adapter ports. The DS8884F includes a base frame (model 984) and a 40U capacity.

- The DS8886 (machine type 283x models 985 and 85E) has a single-phase power storage system and is scalable with up to 24-core processors, up to 192 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives, up to 1,536 standard drives, up to 2048 GB system memory, and up to 128 host adapter ports. The DS8886 includes a base frame (model 985), up to four expansion frames (model 85E), and an expandable 40-46U capacity in each frame.
- The DS8886F (machine type 533x models 985 and 85E) has a single-phase power storage system and is scalable with up to 24-core processors, up to 384 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives, up to 2048 GB system memory, and up to 128 host adapter ports. The DS8886F includes a base frame (model 985) and one expansion frame (model 85E).
- The DS8886 (machine type 283x models 986 and 86E) has a three-phase power storage system and is scalable with up to 24-core processors, up to 192 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives, up to 1,440 standard drives, up to 2048 GB system memory, and up to 128 host adapter ports. The DS8886 includes a base frame (model 986), up to four expansion frames (model 86E), and an expandable 40-46U capacity in each frame.
- The DS8886F (machine type 533x models 986 and 86E) has a three-phase power storage system and is scalable with up to 24-core processors, up to 384 Flash Tier 0, Flash Tier1, or Flash Tier 2 drives, up to 2048 GB system memory, and up to 128 host adapter ports. The DS8886F includes a base frame (model 986) and one expansion frame (model 86E).
- The DS8888F (machine type 533x models 988 and 88E) has a three-phase power storage system and is scalable with up to 48-core processors, up to 768 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives, up to 2048 GB system memory, and up to 128 host adapter ports. The DS8888F includes a base frame (model 988) and up to two expansion frame (model 88E).

The DS8880 features standard 19-inch wide frames and 19-inch wide frames with 6U extensions (DS8886 and DS8888F only).

Energy efficient DC-UPS modules provide 8 KW of single-phase power (model 984 and 985) and 8 KW of three-phase power (model 986 and 988).

DS8880 integrates High Performance Flash Enclosures Gen2 and flash drives for all models documented here to provide a higher level of performance. Previously available models (980, 98B, 981, 98E, 982, 98F) integrated High-Performance Flash Enclosures. For information on models 980, 98B, 981, 98E, 982, and 98F, refer to previous version documentation (GC27-8525-06).

Licensed functions are available in four groups:

Base Function

The Base Function license is required for each DS8880 storage system. The licensed functions include Database Protection, Encryption Authorization, Easy Tier, I/O Priority Manager, the Operating Environment License, and Thin Provisioning.

z-synergy Services

The z-synergy Services include z/OS® functions that are supported on the storage system. The licensed functions include zHyperLink, transparent cloud tiering, High Performance FICON® for z Systems®, HyperPAV, PAV, and z/OS Distributed Data Backup.

Copy Services

Copy Services features help you implement storage solutions to keep your business running 24 hours a day, 7 days a week by providing data duplication, data migration, and disaster recovery functions. The licensed functions include Global Mirror, Metro Mirror, Metro/Global Mirror, Point-in-Time Copy/FlashCopy®, z/OS Global Mirror, Safeguarded Copy, and z/OS Metro/Global Mirror Incremental Resync (RMZ).

Copy Services Manager on Hardware Management Console

The Copy Services Manager on Hardware Management Console (CSM on HMC) license enables IBM Copy Services Manager to run on the Hardware Management Console, which eliminates the need to maintain a separate server for Copy Services functions.

DS8880 also includes features such as:

- POWER8® 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
- 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

DS8880 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 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.

Machine types overview

There are several machine type options available for the DS8000 series. 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 modular rack-mountable system DS8882F model 983 is not documented in this publication. For information on the DS8882F model 983, refer to the *DS8882F Introduction and Planning Guide* (GC27-9259-00).

The following tables list the available hardware machine types and their corresponding function authorization machine types.

Table 4. Available hardware and function-authorization machine types that support both standard drive enclosures and High Performance Flash Enclosures Gen2

Hardware		Licensed functions	
Hardware machine type	Available hardware models that support both standard drive enclosures and High Performance Flash Enclosures Gen2	Corresponding function authorization machine type	Available function authorization models
2831 (1-year warranty period)	984 and 84E 985 and 85E 986 and 86E	2836 (1-year warranty period)	LF8
2832 (2-year warranty period)		2837 (2-year warranty period)	
2833 (3-year warranty period)		2838 (3-year warranty period)	
2834 (4-year warranty period)		2839 (4-year warranty period)	

Table 5. Available hardware and function-authorization machine types that support all-flash system types

Hardware		Licensed functions	
Hardware machine type	Available hardware models that support High Performance Flash Enclosures Gen2	Corresponding function authorization machine type	Available function authorization models
5331 (1-year warranty period)	984 985 and 85E 986 and 86E 988 and 88E	9046 (1-year warranty period)	LF8
5332 (2-year warranty period)		9047 (2-year warranty period)	
5333 (3-year warranty period)		9048 (3-year warranty period)	
5334 (4-year warranty period)		9049 (4-year warranty period)	

Note: Previously available DS8880 models (980, 98B, 981, 98E, 982, 98F) are still supported, but not covered in this version documentation. For information on models not documented here, refer to previous version documentation (GC27-8525-06).

The machine types for the DS8000 series specify the service warranty period. The warranty is used for service entitlement checking when notifications for service are called home. All DS8000 series models report 2107 as the machine type to attached host systems.

Hardware

The architecture of the IBM DS8000 series 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 series architecture

The DS8000 series 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

System types

Starting with version 8.2.1, DS8880 supports five system types with frames: DS8884 (machine type 283x models 984 and 84E), DS8884F (machine type 533x model 984), DS8886 (machine type 283x single-phase power models 985 and 85E, or three-phase power models 986 and 86E), DS8886F (machine type 533x single-phase power models 985 and 85E, or three-phase power models 986 and 86E), and DS8888F (machine type 533x models 988 and 88E).

Note: The modular rack-mountable system DS8882F model 983 is covered in a separate publication. For information on the DS8882F model 983, refer to the *DS888F Introduction and Planning Guide* (GC27-9259-00).

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

DS8884 (machine type 283x models 984 and 84E)

The DS8884 is an entry-level, high-performance, high-capacity storage system that includes standard disk enclosures and High Performance Flash Enclosures Gen2.

DS8884 storage systems feature 6-core (12-core with zHyperLink support) processors and are scalable and support up to 96 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives, and up to 768 standard drives. They are optimized and configured for cost, by minimizing the number of device adapters and maximizing the number of storage enclosures that are attached to each storage system. The frame is 19 inches wide and 40U high.

Model 984 supports the following storage enclosures:

- Up to 4 standard drive enclosure pairs and up to 1 High Performance Flash Enclosure Gen2 pair in a base frame (model 984).
- Up to 5 standard drive enclosure pairs and up to 1 High Performance Flash Enclosure Gen2 pair in a first expansion frame (model 84E).
- Up to 7 standard drive enclosure pairs in a second expansion frame (model 84E).

The DS8884 uses 8 or 16 Gbps Fibre Channel host adapters that run Fibre Channel Protocol (FCP), FICON, or Fibre Channel Arbitrated Loop (FC-AL) (for 8 Gbps adapters only) protocol. The High Performance FICON (HPF) feature is also supported.

The DS8884 supports single-phase power.

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 DS8884, depending on the amount of memory that is available.

Table 6. Components for the DS8884 (models 984 and 84E)

Processors	System memory	Processor memory	I/O enclosure pairs	Host adapters (8 or 4 port)	Device adapter pairs	Standard drive enclosure pairs ^{1, 2}	High Performance Flash Enclosure Gen2 pairs ^{1, 2, 3}	Expansion frames
6-core	64 GB	32 GB	1	2 - 8	0 - 1	0 - 4	0 - 1	0
6-core	128 GB	64 GB	2	2 - 16	0 - 4	0 - 16	0 - 2	0 - 2
6-core	256 GB	128 GB	2	2 - 16	0 - 4	0 - 16	0 - 2	0 - 2
12-core ⁴	256 GB	128 GB	2	2 - 16	0 - 4	0 - 16	0 - 2	0 - 2

1. Standard drive and High Performance Flash Enclosures Gen2 are installed in pairs.
2. This configuration of the DS8880 must be populated with either one standard drive enclosure pair (feature code 1241) or one High Performance Flash Enclosure Gen2 pair (feature code 1600).
3. Each High Performance Flash Enclosure Gen2 pair (feature code 1600) includes a pair of flash RAID adapters.
4. 12-core processors are supported only with zHyperLink.

Table 7. Maximum capacity for the DS8884 (models 984 and 84E)

Processors	System memory	Maximum 2.5-in. standard disk drives	Maximum storage capacity for 2.5-in. standard disk drives	Maximum 3.5-in. standard disk drives	Maximum storage capacity for 3.5-in. standard disk drives	Maximum 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum storage capacity for 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum total drives ¹
6-core	64 GB	192	346 TB	96	576 TB	48	365 TB	240
6-core	128 GB	768	1.38 PB	384	2.3 PB	96	730 TB	864
6-core	256 GB	768	1.38 PB	384	2.3 PB	96	730 TB	864
12-core ²	256 GB	768	1.38 PB	384	2.3 PB	96	730 TB	864

1. Combined total of 2.5-in. standard disk drives and 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives.
2. 12-core processors are supported only with zHyperLink.

Base frame (model 984) overview:

The DS8884 (model 984) includes a base frame.

The base frame includes the following components:

- Standard storage enclosures
- High Performance Flash Enclosures Gen2
- Hardware Management Console (HMC), including a keyboard and monitor
- Second HMC (optional)
- Ethernet switches
- Processor nodes (available with POWER8 processors)
- I/O enclosures
- Direct-current uninterruptible power supplies (DC-UPS)
- Rack power control (RPC) cards

Expansion frame (model 84E) overview:

The DS8884 (model 84E) supports up to two expansion frames that can be added to a base frame. A minimum of 128 GB system memory is required, if expansion frames are added.

The first expansion frame supports up to 240 standard 2.5-inch disk drives. The second expansion frame supports up to 336 2.5-inch standard disk drives. When all three frames are installed, the DS8884 (models 984 and 84E) can support a total of 768 2.5-inch standard disk drives in a compact footprint, creating a high-density storage system and preserving valuable floor space in data center environments.

Only the first expansion frame includes I/O enclosures. You can add up to one High Performance Flash Enclosure Gen2 pair to the first expansion frame. The second expansion frame does not include I/O enclosures or High Performance Flash Enclosures Gen2.

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.

DS8884 (models 984 and 84E) expansion frame location options:

In addition to the standard expansion frame location, the DS8884 offers a remote expansion frame option.

With the standard DS8884 expansion frame location, the first expansion frame is located next to the base frame, and the second expansion frame is located next to the first expansion frame.




Figure 2. DS8884 standard expansion frame locations

With the DS8884 remote expansion frame option, the first expansion frame is located next to the base frame, and the second expansion frame can be located up to 20 meters away from the first expansion frame. This option requires the extended drive cable group C (feature code 1266).




Figure 3. DS8884 remote expansion frame option

DS8884F (machine type 533x model 984)

The DS8884F is an entry-level, high-performance, high-capacity storage system that includes only High Performance Flash Enclosures Gen2.

DS8884F storage systems feature 6-core (12-core with zHyperLink support) processors and are scalable and support up to 192 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives. The frame is 19 inches wide and 40U high.

DS8884F supports up to four High Performance Flash Enclosure Gen2 pair in a base frame (model 984).

The DS8884F uses 8 or 16 Gbps Fibre Channel host adapters that run Fibre Channel Protocol (FCP), FICON, or Fibre Channel Arbitrated Loop (FC-AL) (for 8 Gbps adapters only) protocol. The High Performance FICON (HPF) feature is also supported.

The DS8884F supports single-phase power.

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 DS8884F, depending on the amount of memory that is available.

Table 8. Components for the DS8884F (model 984)

Processors	System memory	Processor memory	I/O enclosure pairs	Host adapters (8 or 4 port)	High Performance Flash Enclosure Gen2 pairs ^{1, 2, 3}	Expansion frames
6-core	64 GB	32 GB	1	2 - 8	1	N/A
6-core	128 GB	64 GB	1 - 2	2 - 16	1 - 4	N/A
6-core	256 GB	128 GB	1 - 2	2 - 16	1 - 4	N/A
12-core ⁴	256 GB	128 GB	1 - 2	2 - 16	1 - 4	N/A
1. High Performance Flash Enclosures Gen2 are installed in pairs. 2. This configuration of the DS8880 must be populated with at least one High Performance Flash Enclosure Gen2 pair (feature code 1600). 3. Each High Performance Flash Enclosure Gen2 pair (feature code 1600) includes a pair of flash RAID adapters. 4. 12-core processors are supported only with zHyperLink.						

Table 9. Maximum capacity for the DS8884F (model 984)

Processors	System memory	Maximum 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum storage capacity for 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum total drives
6-core	64 GB	96	730 TB	96
6-core	128 GB	192	1459 TB	192
6-core	256 GB	192	1459 TB	192
12-core	256 GB	192	1459 TB	192

DS8884F base frame (model 984) overview:

The DS8884F (model 984) includes a base frame.

The base frame includes the following components:

- High Performance Flash Enclosures Gen2
- Hardware Management Console (HMC), including a keyboard and monitor
- Second HMC (optional)
- Ethernet switches
- Processor nodes (available with POWER8 processors)

- I/O enclosure (optional second I/O enclosure with 128 GB or more system memory)
- Direct-current uninterruptible power supplies (DC-UPS)
- Rack power control (RPC) cards

DS8886 (machine type 283x models 985 and 85E or 986 and 86E)

The DS8886 is a high-density, high-performance, high-capacity storage system that includes standard disk enclosures and High Performance Flash Enclosures Gen2.

The DS8886 models 985 and 85E support single-phase power. The DS8886 models 986 and 86E support three-phase power.

DS8886 (machine type 283x models 985 and 85E):

The DS8886 (machine type 283x models 985 and 85E) is a high-density, high-performance, high-capacity storage system that includes standard disk enclosures and High Performance Flash Enclosures Gen2, and supports single-phase power.

DS8886 (machine type 283x models 985 and 85E) storage systems are scalable with up to 24-core processors, up to 192 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives, 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. The frame is 19 inches wide and expandable from 40U - 46U. They support the following storage enclosures:

- Up to 3 standard drive enclosure pairs and up to 2 High Performance Flash Enclosure Gen2 pairs in a base frame (model 985).
- Up to 5 standard drive enclosure pairs and up to 2 High Performance Flash Enclosure Gen2 pairs in a first expansion frame (model 85E).
- Up to 9 standard drive enclosure pairs in a second expansion frame.
- Up to 9 standard drive enclosure pairs in a third expansion frame.
- Up to 6 standard drive enclosure pairs in a fourth expansion frame.

The DS8886 uses 8 or 16 Gbps Fibre Channel host adapters that run Fibre Channel Protocol (FCP), FICON, or Fibre Channel Arbitrated Loop (FC-AL) (for 8 Gbps adapters only) protocol. The High Performance FICON (HPF) feature is also supported.

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 DS8886 (models 985 and 85E), depending on the amount of memory that is available.

Table 10. Components for the DS8886 (machine type 283x models 985 and 85E)

Processors	System memory	Processor memory	I/O enclosure pairs	Host adapters (8 or 4 port)	Device adapter pairs	Standard drive enclosure pairs ¹ ₂	High Performance Flash Enclosure Gen2 pairs ^{1, 2, 3}	Expansion frames
8-core	128 GB 256 GB	64 GB 128 GB	2	2 - 16	0 - 3	0 - 3	0 - 2	0
16-core	256 GB 512 GB	128 GB 256 GB	4	2 - 32	0 - 8	0 - 32	0 - 4	0 - 4
24-core	1024 GB 2048 GB	512 GB 1024 GB	4	2 - 32	0 - 8	0 - 32	0 - 4	0 - 4

Table 10. Components for the DS8886 (machine type 283x models 985 and 85E) (continued)

Processors	System memory	Processor memory	I/O enclosure pairs	Host adapters (8 or 4 port)	Device adapter pairs	Standard drive enclosure pairs ^{1, 2}	High Performance Flash Enclosure Gen2 pairs ^{1, 2, 3}	Expansion frames
- Standard drive and High Performance Flash Enclosures Gen2 are installed in pairs. - This configuration of the DS8886 must be populated with either one standard drive enclosure pair (feature code 1241) or one High Performance Flash Enclosure Gen2 pair (feature code 1600). - Each High Performance Flash Enclosure Gen2 pair (feature code 1600) includes a pair of flash RAID adapters.								

Table 11. Maximum capacity for the DS8886 (machine type 283x models 985 and 85E)

Processors	System memory	Maximum 2.5-in. standard disk drives	Maximum storage capacity for 2.5-in. standard disk drives	Maximum 3.5-in. standard disk drives	Maximum storage capacity for 3.5-in. standard disk drives	Maximum 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum storage capacity for 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum total drives ¹
8-core	128 GB 256 GB	144	259 TB	72	432 TB	96	730 TB	240
16-core	256 GB 512 GB	1536	2.76 PB	768	4.61 PB	192	1459 TB	1728
24-core	1024 GB 2048 GB	1536	2.76 PB	768	4.61 PB	192	1459 TB	1728
1. Combined total of 2.5-in. disk drives and 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives.								

Base frame (model 985) overview:

The DS8886 (model 985) includes a base frame.

The base frame includes the following components:

- Standard storage enclosures
- High Performance Flash Enclosures Gen2
- Hardware Management Console (HMC), including a keyboard and monitor
- Second HMC (optional)
- Ethernet switches
- Processor nodes (available with POWER8 processors)
- I/O enclosures
- Direct-current uninterruptible power supplies (DC-UPS)
- Rack power control (RPC) cards

Expansion frame (model 85E) overview:

The DS8886 supports up to four expansion frames (model 85E) that can be added to a base frame (model 985). A minimum of 256 GB system memory and a 16-core processor is required, if expansion frames are added.

The first expansion frame supports up to 240 2.5-inch standard disk drives. The second and third expansion frames support up to 432 2.5-inch standard disk drives. A fourth expansion frame supports an extra 288 2.5-inch standard disk drives. When all four frames are added, the DS8886 can support a total of 1,536 2.5-inch disk drives in a compact footprint, creating a high-density storage system and preserving valuable floor space in data center environments.

Only the first expansion frame includes I/O enclosures. You can add up to two High Performance Flash Enclosure Gen2 pairs to the first expansion frame. The second, third, and fourth expansion frames do not include I/O enclosures or High Performance Flash Enclosures Gen2.

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.

DS8886 (machine type 283x models 986 and 86E):

The DS8886 (machine type 283x models 986 and 86E) is a high-density, high-performance, high-capacity storage system that includes standard disk enclosures and High Performance Flash Enclosures Gen2, and supports three-phase power.

DS8886 (machine type 283x models 986 and 86E) storage systems are scalable with up to 24-core processors, up to 192 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives, and up to 1,440 standard drives. They are optimized and configured for performance and throughput, by maximizing the number of device adapters and paths to the storage enclosures. The frame is 19 inches wide and expandable from 40U - 46U. They support the following storage enclosures:

- Up to 2 standard drive enclosure pairs and up to 2 High Performance Flash Enclosure Gen2 pairs in a base frame (model 986).
- Up to 4 standard drive enclosure pairs and up to 2 High Performance Flash Enclosure Gen2 pairs in a first expansion frame (model 86E).
- Up to 9 standard drive enclosure pairs in a second expansion frame.
- Up to 9 standard drive enclosure pairs in a third expansion frame.
- Up to 9 standard drive enclosure pairs in a fourth expansion frame.

The DS8886 uses 8 or 16 Gbps Fibre Channel host adapters that run Fibre Channel Protocol (FCP), FICON, or Fibre Channel Arbitrated Loop (FC-AL) (for 8 Gbps adapters only) protocol. The High Performance FICON (HPF) feature is also supported.

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 DS8886 (models 986 and 86E), depending on the amount of memory that is available.

Table 12. Components for the DS8886 (machine type 283x models 986 and 86E)

Processors	System memory	Processor memory	I/O enclosure pairs	Host adapters (8 or 4 port)	Device adapter pairs	Standard drive enclosure pairs ¹ ₂	High Performance Flash Enclosure Gen2 pairs ^{1, 2, 3}	Expansion frames
8-core	128 GB 256 GB	64 GB 128 GB	2	2 - 16	0 - 2	0 - 2	0 - 2	0
16-core	256 GB 512 GB	128 GB 256 GB	4	2 - 32	0 - 8	0 - 30	0 - 4	0 - 4
24-core	1024 GB 2048 GB	512 GB 1024 GB	4	2 - 32	0 - 8	0 - 30	0 - 4	0 - 4

Table 12. Components for the DS8886 (machine type 283x models 986 and 86E) (continued)

Processors	System memory	Processor memory	I/O enclosure pairs	Host adapters (8 or 4 port)	Device adapter pairs	Standard drive enclosure pairs ^{1, 2}	High Performance Flash Enclosure Gen2 pairs ^{1, 2, 3}	Expansion frames
- Standard drive and High Performance Flash Enclosures Gen2 are installed in pairs. - This configuration of the DS8886 must be populated with either one standard drive enclosure pair (feature code 1241) or one High Performance Flash Enclosure Gen2 pair (feature code 1600). - Each High Performance Flash Enclosure Gen2 pair (feature code 1600) includes a pair of flash RAID adapters.								

Table 13. Maximum capacity for the DS8886 (machine type 283x models 986 and 86E)

Processors	System memory	Maximum 2.5-in. standard disk drives	Maximum storage capacity for 2.5-in. standard disk drives	Maximum 3.5-in. standard disk drives	Maximum storage capacity for 3.5-in. standard disk drives	Maximum 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum storage capacity for 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum total drives ¹
8-core	128 GB 256 GB	96	173 TB	48	288 TB	96	730 TB	192
16-core	256 GB 512 GB	1440	2.59 PB	720	4.32 PB	192	1459 TB	1632
24-core	1024 GB 2048 GB	1440	2.59 PB	720	4.32 PB	192	1459 TB	1632
1. Combined total of 2.5-in. disk drives and 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives.								

Base frame (model 986) overview:

The DS8886 (model 986) includes a base frame.

The base frame includes the following components:

- Standard storage enclosures
- High Performance Flash Enclosures Gen2
- Hardware Management Console (HMC), including a keyboard and monitor
- Second HMC (optional)
- Ethernet switches
- Processor nodes (available with POWER8 processors)
- I/O enclosures
- Direct-current uninterruptible power supplies (DC-UPS)
- Rack power control (RPC) cards

Expansion frame (model 86E) overview:

The DS8886 supports up to four expansion frames (model 86E) that can be added to a base frame (model 986). A minimum of 256 GB system memory and a 16-core processor is required, if expansion frames are added.

The first expansion frame supports up to 192 2.5-inch standard disk drives. The second, third, and fourth expansion frames support up to 384 2.5-inch standard disk drives. When all four frames are used, the DS8886 can support a total of 1,440 2.5-inch standard disk drives in a compact footprint, creating a high-density storage system and preserving valuable floor space in data center environments.

Only the first expansion frame includes I/O enclosures. You can add up to two High Performance Flash Enclosure Gen2 pairs to the first expansion frame. The second, third, and fourth expansion frames do not include I/O enclosures or High Performance Flash Enclosures Gen2.

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.

DS8886 expansion frame location options:

In addition to the standard expansion frame location, the DS8886 offers four remote expansion frame options that allow expansion frames to be located up to 20 meters apart.

With the standard DS8886 expansion frame location, the first expansion frame is located next to the base frame, the second expansion frame is located next to the first expansion frame, and each consecutive expansion frame is located next to the previous one.




Figure 4. DS8886 standard expansion frame locations

The DS8886 offers a remote expansion frame option with one remote expansion frame. This option requires the extended drive cable group E (feature code 1254).




Figure 5. DS8886 with one remote expansion frame

The DS8886 offers a remote expansion frame option with two remote expansion frames. This option requires the extended drive cable group D (feature code 1253).




Figure 6. DS8886 with two remote expansion frames

The DS8886 offers a remote expansion frame option with three remote expansion frames. This option requires the extended drive cable group C (feature code 1252).




Figure 7. DS8886 with three remote expansion frames

The DS8886 offers a remote expansion frame option with three separate remote expansion frames. This option requires the extended drive cable groups C, D, and E (feature codes 1252, 1253, and 1254).




Figure 8. DS8886 with three separate remote expansion frames

DS8886F (machine type 533x models 985 and 85E or 986 and 86E)

The DS8886F is a high-density, high-performance, high-capacity storage system that includes only High Performance Flash Enclosures Gen2.

The DS8886F models 985 and 85E support single-phase power. The DS8886F models 986 and 86E support three-phase power.

DS8886F (machine type 533x models 985 and 85E):

The DS8886F (machine type 533x models 985 and 85E) is a high-density, high-performance, high-capacity storage system that includes High Performance Flash Enclosures Gen2, and supports single-phase power.

DS8886F (machine type 533x models 985 and 85E) storage systems are scalable with up to 24-core processors, and up to 384 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives. They are optimized and configured for performance and throughput by maximizing the number of paths to the storage enclosures. The frame is 19 inches wide and 46U high. They support the following storage enclosures:

- Up to 4 High Performance Flash Enclosure Gen2 pairs in a base frame (model 985).
- Up to 4 High Performance Flash Enclosure Gen2 pairs in an expansion frame (model 85E).

The DS8886F uses 8 or 16 Gbps Fibre Channel host adapters that run Fibre Channel Protocol (FCP), FICON, or Fibre Channel Arbitrated Loop (FC-AL) (for 8 Gbps adapters only) protocol. The High Performance FICON (HPF) feature is also supported.

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 DS8886F (models 985 and 85E), depending on the amount of memory that is available.

Table 14. Components for the DS8886F (machine type 533x models 985 and 85E)

Processors	System memory	Processor memory	I/O enclosure pairs	Host adapters (8 or 4 port)	High Performance Flash Enclosure Gen2 pairs ^{1, 2, 3}	Expansion frames
8-core	128 GB 256 GB	64 GB 128 GB	2	2 - 16	1 - 4	0
16-core	256 GB 512 GB	128 GB 256 GB	4	2 - 32	1 - 8	0 - 1
24-core	1024 GB 2048 GB	512 GB 1024 GB	4	2 - 32	1 - 8	0 - 1
1. High Performance Flash Enclosures Gen2 are installed in pairs. 2. This configuration of the DS8880 must be populated with at least one High Performance Flash Enclosure Gen2 pair (feature code 1600). 3. Each High Performance Flash Enclosure Gen2 pair (feature code 1600) includes a pair of flash RAID adapters.						

Table 15. Maximum capacity for the DS8886F (machine type 533x models 985 and 85E)

Processors	System memory	Maximum 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum storage capacity for 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum total drives
8-core	128 GB 256 GB	192	1459 TB	192
16-core	256 GB 512 GB	384	2918 TB	384
24-core	1024 GB 2048 GB	384	2918 TB	384

DS8886F base frame (model 985) overview:

The DS8886F (model 985) includes a base frame.

The base frame includes the following components:

- High Performance Flash Enclosures Gen2
- Hardware Management Console (HMC), including a keyboard and monitor
- Second HMC (optional)
- Ethernet switches
- Processor nodes (available with POWER8 processors)
- I/O enclosures
- Direct-current uninterruptible power supplies (DC-UPS)
- Rack power control (RPC) cards

DS8886F expansion frame (model 85E) overview:

The DS8886F supports one expansion frame (model 85E) that can be added to a base frame (model 985). A minimum of 256 GB system memory and a 16-core processor is required to add the expansion frame.

The expansion frame includes I/O enclosures. You can add up to four High Performance Flash Enclosure Gen2 pairs to the expansion frame.

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

DS8886F (machine type 533x models 986 and 86E):

The DS8886F (machine type 533x models 986 and 86E) is a high-density, high-performance, high-capacity storage system that includes High Performance Flash Enclosures Gen2, and supports three-phase power.

DS8886F (machine type 533x models 986 and 86E) storage systems are scalable with up to 24-core processors, and up to 384 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives. They are optimized and configured for performance and throughput by maximizing the number of paths to the storage enclosures. The frame is 19 inches wide and 46U high. They support the following storage enclosures:

- Up to 4 High Performance Flash Enclosure Gen2 pairs in a base frame (model 986).
- Up to 4 High Performance Flash Enclosure Gen2 pairs in an expansion frame (model 86E).

The DS8886F uses 8 or 16 Gbps Fibre Channel host adapters that run Fibre Channel Protocol (FCP), FICON, or Fibre Channel Arbitrated Loop (FC-AL) (for 8 Gbps adapters only) protocol. The High Performance FICON (HPF) feature is also supported.

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 DS8886F (models 986 and 86E), depending on the amount of memory that is available.

Table 16. Components for the DS8886F (machine type 533x models 986 and 86E)

Processors	System memory	Processor memory	I/O enclosure pairs	Host adapters (8 or 4 port)	High Performance Flash Enclosure Gen2 pairs ^{1, 2, 3}	Expansion frames
8-core	128 GB 256 GB	64 GB 128 GB	2	2 - 16	1 - 4	0
16-core	256 GB 512 GB	128 GB 256 GB	4	2 - 32	1 - 8	0 - 1
24-core	1024 GB 2048 GB	512 GB 1024 GB	4	2 - 32	1 - 8	0 - 1

1. High Performance Flash Enclosures Gen2 are installed in pairs.

2. This configuration of the DS8880 must be populated with at least one High Performance Flash Enclosure Gen2 pair (feature code 1600).

3. Each High Performance Flash Enclosure Gen2 pair (feature code 1600) includes a pair of flash RAID adapters.

Table 17. Maximum capacity for the DS8886F (machine type 533x models 986 and 86E)

Processors	System memory	Maximum 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum storage capacity for 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum total drives
8-core	128 GB 256 GB	192	1459 TB	192

Table 17. Maximum capacity for the DS8886F (machine type 533x models 986 and 86E) (continued)

Processors	System memory	Maximum 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum storage capacity for 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum total drives
16-core	256 GB 512 GB	384	2918 TB	384
24-core	1024 GB 2048 GB	384	2918 TB	384

DS8886F base frame (model 986) overview:

The DS8886F (model 986) includes a base frame.

The base frame includes the following components:

- High Performance Flash Enclosures Gen2
- Hardware Management Console (HMC), including a keyboard and monitor
- Second HMC (optional)
- Ethernet switches
- Processor nodes (available with POWER8 processors)
- I/O enclosures
- Direct-current uninterruptible power supplies (DC-UPS)
- Rack power control (RPC) cards

DS8886F expansion frame (model 86E) overview:

The DS8886F supports one expansion frame (model 86E) that can be added to a base frame (model 986). A minimum of 256 GB system memory and a 16-core processor is required for the expansion frame to be added.

The expansion frame includes I/O enclosures. You can add up to four High Performance Flash Enclosure Gen2 pairs to the expansion frame.

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

DS8888F (machine type 533x models 988 and 88E)

The DS8888F (machine type 533x models 988 and 88E) is a high-performance, high-efficiency, high-capacity storage system that includes only High Performance Flash Enclosures Gen2.

DS8888F storage systems (machine type 533x models 988 and 88E) are scalable with up to 48-core processors, 16 High Performance Flash Enclosures Gen2 pairs, and up to 768 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives. They are optimized and configured for cost. The frame is 19 inches wide with a 46U capacity. They support the following storage enclosures:

- Up to 4 High Performance Flash Enclosure Gen2 pairs in the base frame (model 988).
- Up to 6 High Performance Flash Enclosure Gen2 pairs each in the first expansion frame (model 88E).
- Up to 6 High Performance Flash Enclosure Gen2 pairs each in the second expansion frame.

The DS8888F uses 8 or 16 Gbps Fibre Channel host adapters that run Fibre Channel Protocol (FCP), FICON, or (for 8 Gbps adapters only) Fibre Channel Arbitrated Loop (FC-AL) protocol. The High Performance FICON (HPF) feature is also supported.

The DS8888F supports three-phase power only.

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 DS8888F (models 988 and 88E), depending on the amount of memory that is available.

Table 18. Components for the DS8888F (machine type 533x models 988 and 88E)

Processors	System memory	Processor memory	I/O enclosures	Host adapters (8 or 4 port)	High Performance Flash Enclosure Gen2 pairs ^{1, 2, 3}	Expansion frames
24-core	1024 GB	512 GB	4	2 - 16	1 - 4	0
48-core	2048 GB	1024 GB	8	2 - 32	1 - 16	0 - 2
- High Performance Flash Enclosures Gen2 are installed in pairs. - A maximum of 8 feature code 1600 are supported in the system. For the remaining 8 High Performance Flash Enclosure Gen2 pair, feature code 1602 should be ordered with feature code 1604.						

Table 19. Maximum capacity for the DS8888F (machine type 533x models 988 and 88E)

Processors	System memory	Maximum 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum storage capacity for 2.5-in. Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives	Maximum total drives
24-core	1024 GB	192	1459 TB	192
48-core	2048 GB	768	5837 TB	768

DS8888F base frame (model 988) overview:

The DS8888F includes a base frame (model 988).

The base frame includes the following components:

- High Performance Flash Enclosures Gen2
- Hardware Management Console (HMC), including a keyboard and monitor
- Second HMC (optional)
- Ethernet switches
- Processor nodes (available with POWER8+ processors)
- I/O enclosures
- Direct-current uninterruptible power supplies (DC-UPS)
- Rack power control (RPC) cards

DS8888F expansion frame (model 88E) overview:

The DS8888F supports up to two expansion frames (model 88E) that can be added to a base frame. A minimum of 2048 GB system memory and a 48-core processor is required, if an expansion frame is used.

The first expansion frame includes I/O enclosures. You can add up to six High Performance Flash Enclosure Gen2 pairs to each expansion frame.

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.

Storage enclosures

DS8880 version 8.2 and later integrates one of two types of storage enclosures: High Performance Flash Enclosures Gen2 and standard drive enclosures.

High Performance Flash Enclosures Gen2 pair

The High Performance Flash Enclosure Gen2 is a 2U storage enclosure that is installed in pairs.

The High Performance Flash Enclosure Gen2 pair provides two 2U storage enclosures with associated RAID controllers and cabling. This combination of components forms a high-performance, fully-redundant flash storage array.

The High Performance Flash Enclosure Gen2 pair contains the following hardware components:

- Two 2U 24-slot SAS flash drive enclosures. Each of the two enclosures contains the following components:
 - Two power supplies with integrated cooling fans
 - Two SAS Expander Modules with two SAS ports each
 - One midplane or backplane for plugging components that provides maintenance of flash drives, Expander Modules, and power supplies
- Two High Performance Flash Enclosure Gen2 flash RAID adapters configured for redundant access to the 2U flash enclosures. Each RAID adapter supports concurrent maintenance and includes the following components:
 - High Performance ASIC RAID engine
 - Four SAS ports and cables connected to the four SAS Expander Modules providing fully-redundant access from each RAID adapter to both 2U enclosures
 - Integrated cooling
 - x8 PCIe Gen2 cable port for direct connection to the I/O enclosure

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 and flash drives (SSDs) are available in sets of 16 drives. The 3.5-inch SAS disk drives are available in half-drive sets of eight drives.

Management console

The management console 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 monitor that are stored on the left side of the base frame.

The HMC connects to the customer network and provides access to 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 storage system.

Each base frame includes one HMC and space for a second HMC, which 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 contains POWER8 processors and memory.

I/O enclosures

The I/O enclosure provides connectivity between the adapters and the processor complex.

The I/O enclosure uses PCIe interfaces to connect 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, the I/O enclosure is connected to each processor node with a point-to-point connection.

The I/O enclosure contain the following adapters:

Flash interface connectors

Interface connector that provides PCIe cable connection from the I/O enclosure to the High Performance Flash Enclosure Gen2.

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 controllers for RAID support.

Host adapters

An I/O enclosure can support up to 16 host ports. For example, if an 8-port adapter is used, then only 2 additional 4-port adapters can be installed, or one more 8-port adapter. If only 8-port adapters are used, then only 2 host adapters can be installed in each I/O enclosure. If only 4-port adapters are used, then 4 host adapters can be installed in each I/O enclosure.

For PCIe-attached adapters with four or eight 8 Gbps Fibre Channel ports, each port can be independently configured to use SCSI/FCP, SCSI/FC-AL, or FICON/zHPF protocols. For PCIe-attached adapters with 4 16 Gbps Fibre Channel ports, each port can be independently configured to use SCSI/FCP or FICON/zHPF protocols. Both longwave and shortwave adapter versions that support different maximum cable lengths are available. 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 (zHPF) protocol can be used by FICON host channels that have zHPF support. The use of zHPF 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. DS8880 uses three-phase and single-phase power.

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 180 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 40 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.

Functional overview

The following list provides an overview of some of the features that are associated with DS8880.

Note: Some storage system functions are not available 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

DS8880 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 DS8880 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 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 that is created in the cryptoperiod starting in 2014. The DS8880 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 DS8880 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 by using the DS8000 Storage Management GUI, DS CLI, and DS Open API. The rotate capacity allocation method (also referred to as rotate volumes) is an alternative allocation method that tends to prefer volumes that are allocated to a single managed array, and is not recommended. The rotate extents 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 that uses 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 by using IBM Copy Services Manager or IBM Spectrum Control. .

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 DS8880 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 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 DS8880 storage system can work in cooperation with IBM Z hosts to provide the following performance enhancement functions.

- Extended Address Volumes
- High Performance FICON for IBM Z
- I/O Priority Manager with z/OS Workload Manager
- Parallel Access Volumes and HyperPAV (also referred to as aliases)
- Quick initialization for IBM Z
- Transparent cloud tiering
- ZHyperLink technology that speeds up transaction processing and improves active log throughput

Copy Services

The DS8880 storage system supports a wide variety of Copy Service

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

- FlashCopy
- Remote Pair FlashCopy (Preserve Mirror)
- Safeguarded Copy
- Remote Mirror and Copy:
 - Metro Mirror
 - Global Copy
 - Global Mirror
 - Metro/Global Mirror
 - Multi-Target PPRC
 - z/OS Global Mirror
 - z/OS Metro/Global Mirror

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, if 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. If I/O Priority Manager is enabled on your storage system, you cannot use a zHyperLink connection.

Restriction of hazardous substances (RoHS)

The DS8880 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. Although the end result of storage configuration is similar, each interface has specific terminology, concepts and procedures.

Note: LSS is synonymous with logical control unit (LCU) and subsystem identification (SSID).

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 9 illustrates the concepts of configuration.




Figure 9. 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 8 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 storage system. Hosts defined on the storage system are configured with a user-designated host type that enables the storage system to recognize and interact with the host. Only hosts that are mapped to volumes can access those volumes.

Logical configuration of the 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 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 IBM Z 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 volumes between 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 IBM Z host, the next steps of the configuration process are completed on the host.

If the volumes are connecting to an open systems host, map the volumes to the host, add host ports to the host, and then map the ports to the 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, a large extent is 1 GiB, and a small extent is 16 MiB. If you work in an IBM Z environment, a large extent is the size of an IBM 3390 Mod 1 disk drive (1113 cylinders), and a small extent is 21 cylinders.

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 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 configuration 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. The 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. Identify different volume groups for different uses or functions (for example, SCSI target, 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 10 on page 31 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. The Quick Initialization feature for open system on FB ESE volumes allows quicker access to logical volumes. The volumes include host volumes and source volumes that can be used 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 specify LUN ID numbers through the graphical user interface (GUI) for volumes in a map-type volume group. You can 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 might 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 are expanded from the previous maximum of 16.




Figure 10. Logical configuration sequence

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 some 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.

RAID 6 is the recommended and default RAID type for all drives over 1 TB. RAID 6 and RAID 10 are the only supported RAID types for 3.8 TB Flash Tier 1 drives. RAID 6 is the only supported RAID type for 7.6 TB Flash Tier 2 drives.

RAID 5 overview

RAID 5 is a method of spreading volume data across multiple drives.

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 5 is not supported for drives larger than 1 TB and requires a request for price quote (RPQ). For information, contact your sales representative.

Note: RAID 6 is the recommended and default RAID type for all drives over 1 TB. RAID 6 and RAID 10 are the only supported RAID types for 3.8 TB Flash Tier 1 drives. RAID 6 is the only supported RAID type for 7.6 TB Flash Tier 2 drives.

RAID 6 overview

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

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.

The default RAID type for all drives over 1 TB is RAID 6. RAID 6 and RAID 10 are the only supported RAID types for 3.8 TB Flash Tier 1 drives. RAID 6 is the only supported RAID type for 7.6 TB Flash Tier 2 drives.

RAID 10 overview

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

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.

Note: RAID 6 is the recommended and default RAID type for all drives over 1 TB. RAID 6 and RAID 10 are the only supported RAID types for 3.8 TB Flash Tier 1 drives. RAID 6 is the only supported RAID type for 7.6 TB Flash Tier 2 drives.

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 storage system can contain up to 255 LSSs and each LSS can be connected to 16 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 IBM Z 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 64K 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 4K 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 IBM Z 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.

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 storage system.

These interfaces include DS8000 Storage Management GUI, DS Command-Line Interface (DS CLI), the DS Open Application Programming Interface, DS8000 RESTful API, IBM Storage Mobile Dashboard, IBM Spectrum Control and IBM Copy Services Manager.

DS8000 Storage Management GUI

Use the DS8000 Storage Management GUI to configure and manage storage and monitor performance and Copy Services functions.

DS8000 Storage Management GUI is a web-based GUI that is installed on the Hardware Management Console (HMC). 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 38.

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 the DS8000 Storage Management GUI. For more information about security requirements, see information about configuring your environment for NIST SP 800-131A compliance in the IBM DS8000 series online product documentation (http://www.ibm.com/support/knowledgecenter/ST5GLJ_8.1.0/com.ibm.storage.ssic.help.doc/f2c_securitybp.html).
- User names and passwords are encrypted for HTTPS protocol. You cannot access the DS8000 Storage Management GUI over the non-secure HTTP protocol (port 8451).

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 system. 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 system, including managing security settings, querying point-in-time performance information or status of physical resources, and exporting audit logs.

Note: Java™ 1.8 must be installed on systems that run the DS CLI.

The DS CLI provides a full-function set of commands to manage logical configurations and Copy Services configurations. The DS CLI is available in the

DS8000 Storage Management GUI. The DS CLI client can also be installed on and is supported in many different environments, including the following platforms:

- AIX 6.1, 7.1, 7.2
- Linux, Red Hat Enterprise Linux [RHEL] 6 and 7
- Linux, SUSE Linux, Enterprise Server (SLES) 11 and 12
- VMware ESX 5.5, 6 Console
- IBM i 7.1, 7.2
- Oracle Solaris 10 and 11
- Microsoft Windows Server 2008, 2012 and Windows 7, 8, 8.1, 10

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 client. For more information about security requirements, see documentation about configuring your environment for NIST SP 800-131A compliance in IBM Knowledge Center (https://www.ibm.com/support/knowledgecenter/ST5GLJ_8.5.0/com.ibm.storage.ssic.help.doc/f2c_securitybp_nist.html).

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 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 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.

RESTful API

The RESTful API is an application on the DS8000 HMC for initiating simple storage operations through the Web.

The RESTful (Representational State Transfer) API is a platform independent means by which to initiate create, read, update, and delete operations in the storage system and supporting storage devices. These operations are initiated with the HTTP commands: **POST**, **GET**, **PUT**, and **DELETE**.

The RESTful API is intended for use in the development, testing, and debugging of DS8000 client management infrastructures. You can use the RESTful API with a

CURL command or through standard Web browsers. For instance, you can use the storage system with the RESTClient add-on.

IBM Spectrum Control

IBM Spectrum Control is an integrated software solution that can help you improve and centralize the management of your storage environment through the integration of products. With IBM Spectrum Control, it is possible to manage multiple DS8000 systems from a single point of control.

Note: IBM Spectrum Control is not required for the operation of a storage system. However, it is recommended. IBM Spectrum Control can be ordered and installed as a software product on various servers and operating systems. When you install IBM Spectrum Control, ensure that the selected version supports the current system functions. Optionally, you can order a server on which IBM Spectrum Control is preinstalled.

IBM Spectrum Control simplifies storage management by providing the following benefits:

- Centralizing the management of heterogeneous storage network resources with IBMstorage management software
- Providing greater synergy between storage management software and IBMstorage 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

For more information, see IBM Spectrum Control online product documentation in IBM Knowledge Center (www.ibm.com/support/knowledgecenter).

IBM Copy Services Manager

IBM Copy Services Manager controls Copy Services in storage environments. Copy Services are features that are used by storage systems, such as DS8000, to configure, manage, and monitor data-copy functions.

IBM Copy Services Manager provides both a graphical interface and command line that you can use for configuring and managing Copy Services functions across storage units. Copy Services include the point-in-time function – IBM FlashCopy and Safeguarded Copy, and the remote mirror and copy functions – Metro Mirror, Global Mirror, and Metro Global Mirror. Copy Services Manager can automate the administration and configuration of these services; and monitor and manage copy sessions.

You can use Copy Services Manager to complete the following data replication tasks and help reduce the downtime of critical applications:

- Plan for replication when you are provisioning storage
- Keep data on multiple related volumes consistent across storage systems for a planned or unplanned outage
- Monitor and track replication operations
- Automate the mapping of source volumes to target volumes

Starting with DS8000 Version 8.1, Copy Services Manager also comes preinstalled on the Hardware Management Console (HMC). Therefore, you can enable the

Copy Services Manager software that is already on the hardware system. Doing so results in less setup time; and eliminates the need to maintain a separate server for Copy Services functions.

You can also use Copy Services Manager to connect to an LDAP repository for remote authentication. For more information, see the DS8000 online product documentation at http://www.ibm.com/support/knowledgecenter/ST5GLJ/ds8000_kcwelcome.html and search for topics that are related to *remote authentication*.

For more information, see the Copy Services Manager online product documentation at http://www.ibm.com/support/knowledgecenter/SSESK4/csm_kcwelcome.html. The "What's new" topic provides details on the features added for each version of Copy Services Manager that can be used by DS8000, including HyperSwap for multi-target sessions, and incremental FlashCopy support.

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 20. Supported web browsers

DS8000 series version	Supported browsers
8.5	Mozilla Firefox 38 Mozilla Firefox Extended Support Release (ESR) 38 Microsoft Internet Explorer 11 Google Chrome 43

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 the 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 IBM Spectrum Control with Internet Explorer, complete the following steps to properly configure the web browser.

1. Disable the **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.

For more information, see your browser documentation.

Chapter 2. Hardware features

Use this information to assist you with planning, ordering, and managing your storage system.

The following table lists feature codes that are used to order hardware features for your system.

Table 21. Feature codes for hardware features

Feature code	Feature	Description
0101	Single-phase input power indicator, 200 - 220 V, 30 A	
0102	Single-phase input power indicator, 220 - 240 V, 30 A	
0170	Top expansion	For models 985 and 85E, 986 and 86E, 988 and 88E, increases frame from 40U to 46U
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.
0403	Non-encryption certification key	Required when the storage system model is shipped to China or Russia.
1050	Battery service modules	Single-phase DC-UPS
1052	Battery service modules	Three-phase DC-UPS
1055	Extended power line disturbance	An optional feature that is used to protect the storage system from a power-line disturbance for up to 40 seconds.
1062	Single-phase power cord, 200 - 240 V, 60 A, 3-pin connector	HBL360C6W, Pin and Sleeve Connector, IEC 60309, 2P3W HBL360R6W, AC Receptacle, IEC 60309, 2P3W
1063	Single-phase power cord, 200 - 240 V, 63 A, no connector	Inline Connector: not applicable Receptacle: not applicable
1086	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 60309, 4P5W HBL530R6V02, AC Receptacle, IEC 60309, 4P5W
1087	Three-phase low voltage (four-wire 3P+G), 200-240V, 30 A, IEC 60309 4-pin customer connector	HBL430C9W, Pin and Sleeve Connector, IEC 60309, 3P4W HBL430R9W, AC Receptacle, IEC 60309, 3P4W

Table 21. Feature codes for hardware features (continued)

Feature code	Feature	Description
1088	Three-phase high voltage (five-wire 3P+N+G), 380-415V, 40 A, no customer connector provided	Inline Connector: not applicable Receptacle: not applicable
1089	Three-phase low voltage (four-wire 3P+G), 200-240V, 60 A, IEC 60309 4-pin customer connector	HBL460C9W, Pin and Sleeve Connector, IEC 60309, 3P4W HBL460R9W, AC Receptacle, IEC 60309, 3P4W
1101	5 ft. ladder	For models 984 and 84E, 985 and 85E, 986 and 86E, 988 and 88E
1102	3 ft. platform ladder	For models 984 and 84E, 985 and 85E, 986 and 86E, 988 and 88E
1103	Rolling step stool	For models 984, 985, 986, and 988
1141	Primary management console	A primary management consoles is required to be installed in the model 984, 985, 986, and 988
1151	Secondary management console	Redundant management console for high availability For model 984, 985, 986, and 988, this feature is optional
1241	Drive enclosure pair total	Admin feature totaling all disk enclosure pairs installed in the model
1242	Standard drive enclosure pair	For 2.5-inch disk drives
1244	Standard drive enclosure pair	For 3.5-inch disk drives
1245	Standard drive enclosure pair	For 400 GB flash drives
1246	Drive cable group A	Connects the disk drives to the device adapters within the same base model 985 or 986
1247	Drive cable group B	Connects the disk drives to the device adapters in the first expansion model 85E or 86E
1248	Drive cable group C	Connects the disk drives from a second expansion model 85E or 86E to the base model 985 or 986 and first expansion model 85E or 86E
1249	Drive cable group D	Connects the disk drives from a third expansion model 85E or 86E to a second expansion model 85E or 86E
1251	Drive cable group E	Connects the disk drives from a fourth expansion model 85E or 86E to a third expansion model 85E or 86E
1252	Extended drive cable group C	For model 85E or 86E, 20 meter cable to connect disk drives from a third expansion model 85E or 86E to a second expansion model 85E or 86E

Table 21. Feature codes for hardware features (continued)

Feature code	Feature	Description
1253	Extended drive cable group D	For model 85E or 86E, 20 meter cable to connect disk drives from a fourth expansion model 85E or 86E to a third expansion model 85E or 86E
1254	Extended drive cable group E	For model 85E or 86E, 20 meter cable to connect disk drives from a fifth expansion model 85E or 86E to a fourth expansion model 85E or 86E
1256	Standard drive enclosure pair	For 800 GB flash drives
1257	Standard drive enclosure pair	For 1.6 TB flash drives
1261	Drive cable group A	Connects the disk drives to the device adapters within the same base model 984
1262	Drive cable group B	Connects the disk drives to the device adapters in the first expansion model 84E
1263	Drive cable group C	Connects the disk drives from a second expansion model 84E to the base model 984 and first expansion model 84E
1266	Extended drive cable group C	For model 84E, 20 meter cable to connect disk drives from a second expansion model 84E to the base model 984 and first expansion model 84E
1303	Gen2 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.
1400	Top-exit bracket for Fibre Channel cable	
1410	Fibre Channel cable	40 m (131 ft), 50 micron OM3 or higher, multimode LC
1411	Fibre Channel cable	31 m (102 ft), 50 micron OM3 or higher, multimode LC/SC
1412	Fibre Channel cable	2 m (6.5 ft), 50 micron OM3 or higher, multimode LC/SC Jumper
1420	Fibre Channel cable	31 m (102 ft), 9 micron OS1 or higher, single mode LC
1421	Fibre Channel cable	31 m (102 ft), 9 micron OS1 or higher, single mode LC/SC
1422	Fibre Channel cable	2 m (6.5 ft), 9 micron OS1 or higher, single mode LC/SC Jumper
1600	High Performance Flash Enclosure Gen2 pair with flash RAID controllers	For flash drives

Table 21. Feature codes for hardware features (continued)

Feature code	Feature	Description
1602	High Performance Flash Enclosure Gen2 pair	For DS8888F (models 988 and 88E) Requires feature code 1604
1604	Flash RAID adapter pair	For DS8888F (models 988 and 88E), required after first 8 feature code 1600 Requires feature code 1602
1610	400 GB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives)
1611	800 GB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives)
1612	1.6 TB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives)
1613	3.2 TB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives)
1623	3.8 TB 2.5-inch Flash Tier 1 drive set	Flash drive set (16 drives) No intermix with Flash Tier 0 or Flash Tier 2 drive sets
1624	7.6 TB 2.5-inch Flash Tier 2 drive set	Flash drive set (16 drives) No intermix with Flash Tier 0 or Flash Tier 1 drive sets
1699	High Performance Flash Enclosure Gen2 filler set	Includes 16 fillers
1761	External SKLM isolated-key appliance	Model AP1 single server configuration
1762	Secondary external SKLM isolated-key appliance	Model AP1 dual server configuration
1884	DS8000 Licensed Machine Code R8.4	Microcode bundle 88.x.xx.x for base models 984, 985, 986, and 988
1885	DS8000 Licensed Machine Code R8.5	Microcode bundle 88.x.xx.x for base model 984, 985, 986, and 988
1906	Earthquake resistance kit	
1984	DS8000 Licensed Machine Code R8.4	Microcode bundle 88.x.xx.x for expansion models 84E, 85E, 86E, and 88E
1985	DS8000 Licensed Machine Code R8.5	Microcode bundle 88.x.xx.x for expansion models 84E, 85E, 86E, and 88E
2997	Disk enclosure filler set	For 3.5-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
3153	Fibre Channel host-adapter	4-port, 8 Gbps shortwave FCP and FICON host adapter PCIe
3157	Fibre Channel host-adapter	8-port, 8 Gbps shortwave FCP and FICON host adapter PCIe
3253	Fibre Channel host-adapter	4-port, 8 Gbps longwave FCP and FICON host adapter PCIe

Table 21. Feature codes for hardware features (continued)

Feature code	Feature	Description
3257	Fibre Channel host-adapter	8-port, 8 Gbps longwave FCP and FICON host adapter PCIe
3353	Fibre Channel host-adapter	4-port, 16 Gbps shortwave FCP and FICON host adapter PCIe
3453	Fibre Channel host-adapter	4-port, 16 Gbps longwave FCP and FICON host adapter PCIe Requires feature code 3065 or 3066
3500	zHyperLink I/O-adapter	Required for feature code 1450 and 1451
3600	Transparent cloud tiering adapter pair for 2U processor complex (optional)	2-port 10 Gbps SFP+ optical/2-port 1 Gbps RJ-45 copper longwave adapter pair for model 984
3601	Transparent cloud tiering adapter pair for 4U processor complex (optional)	2-port 10 Gbps SFP+ optical/2-port 1 Gbps RJ-45 copper longwave adapter pair for models 985, 986, and 988
4233	64 GB system memory	(6-core)
4234	128 GB system memory	(6-core)
4235	256 GB system memory	(6-core and 12-core)
4334	128 GB system memory	(8-core)
4335	256 GB system memory	(8-core and 16-core)
4336	512 GB system memory	(16-core)
4337	1 TB system memory	(24-core)
4338	2 TB system memory	(24-core)
4421	6-core POWER8 processors	Requires feature code 4233, 4234, or 4235
4422	8-core POWER8 processors	Requires feature code 4334, or 4335
4423	16-core POWER8 processors	Requires feature code 4335, or 4336
4424	24-core POWER8 processors	Requires feature code 4337, or 4338
4425	12-core POWER8 processors	For zHyperLink support on model 984 Requires feature code 4235
4497	1 TB system memory	(24-core)
4498	2 TB system memory	(48-core)
4878	24-core POWER8+ processors	Requires feature code 4497
4898	48-core POWER8+ processors	Requires feature code 4498
5308	300 GB 15 K FDE disk-drive set	SAS
5618	600 GB 15 K FDE disk-drive set	SAS
5708	600 GB 10K FDE disk-drive set	SAS
5768	1.2 TB 10K FDE disk-drive set	SAS
5778	1.8 TB 10K FDE disk-drive set	SAS
5868	4 TB 7.2 K FDE disk-drive set	SAS
5878	6 TB 7.2 K FDE disk-drive set	SAS

Table 21. Feature codes for hardware features (continued)

Feature code	Feature	Description
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
6458	3.2 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 management console units in one of the storage units. You can add a second management console for redundancy.

Management console

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

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

To provide continuous availability of your access to the management-console functions, use an additional management console , especially for storage environments that use encryption. Both management consoles share a keyboard and display that are stored on the left side of the base frame.

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 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 power control card.

Storage system

The storage unit contains a storage server and one or more pair of storage enclosures that are packaged in one or more frames 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 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.

Power control card

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

Disk drives and flash drives

The storage system provides you with a choice of drives.

The following drives are available:

- 2.5-inch Flash Tier 0 drives with FDE
 - 400 GB
 - 800 GB
 - 1.6 TB
 - 3.2 TB
- 2.5-inch Flash Tier 1 drives with FDE
 - 3.8 TB
- 2.5-inch Flash Tier 2 drives with FDE
 - 7.6 TB
- 2.5-inch flash drives with FDE
 - 400 GB
 - 800 GB
 - 1.6 TB
- 2.5-inch disk drives with FDE
 - 300 GB, 15 K RPM
 - 600 GB, 15 K RPM
 - 600 GB, 10 K RPM
 - 1.2 TB, 10 K RPM
 - 1.8 TB, 10 K RPM
- 3.5-inch disk drives with FDE
 - 4 TB, 7.2 K RPM
 - 6 TB, 7.2 K RPM

Note: Intermix of Flash Tier 0, Flash Tier 1, and Flash Tier 2 drives is not supported.

Drive maintenance policy

The 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. 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 storage system provides various host attachments so that you can consolidate storage capacity and workloads for open-systems hosts and IBM Z.

The storage system 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 systems with frames. These systems can have a maximum of four host adapters per I/O enclosure including 4-port 16 Gbps adapters, 4- or 8-port 8 Gbps adapters, or a combination of each.

Models 984, 985, 986, and 988

A maximum of 16 ports per I/O enclosure is supported, which provides for a maximum of 128 ports in a system. Eight-port 8 Gbps adapters are allowed only in slots C1 and C4. If an 8-port adapter is present in slot C1, no adapter can be installed in slot C2. If an 8-port adapter is present in slot C4, no adapter can be installed in slot C5.

The following table shows the host adapter plug order.

Table 22. Plug order for 4- and 8-port HA slots for two and four I/O enclosures

I/O enclosures	Slot number					
	C1	C2	C3	C4	C5	C6
For two I/O enclosures						
Top I/O enclosure 1						
Bottom I/O enclosure 3	3	7		1	5	
Top I/O enclosure 2						
Bottom I/O enclosure 4	2	8		4	6	
For four I/O enclosures in a DS8884, DS8886 configuration						
Top I/O enclosure 1	7	15		3	11	
Bottom I/O enclosure 3	5	13		1	9	
Top I/O enclosure 2	4	12		8	16	
Bottom I/O enclosure 4	2	10		6	14	

The following HA-type plug order is used during manufacturing when different types of HA cards are installed.

1. 8-port 8 Gbps longwave host adapters
2. 8-port 8 Gbps shortwave host adapters
3. 4-port 16 Gbps longwave host adapters
4. 4-port 16 Gbps shortwave host adapters
5. 4-port 8 Gbps longwave host adapters
6. 4-port 8 Gbps shortwave host adapters

Open-systems host attachment with Fibre Channel adapters

You can attach a storage system to an open-systems host with Fibre Channel adapters.

The storage system supports SAN speeds of up to 16 Gbps with the current 16 Gbps host adapters, or up to 8 Gbps with the 8 Gbps host adapters. The storage system detects and operates at the greatest available link speed that is shared by both sides of the system.

Fibre Channel technology transfers data between the sources and the users of the information. 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 IBM Z hosts overview

The storage system can be attached to FICON attached IBM Z host operating systems under specified adapter configurations.

Each storage system Fibre Channel adapter has four ports or eight ports, depending on the adapter type. Each port has a unique worldwide port name (WWPN). You can configure the port to operate with the FICON upper-layer protocol.

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 64 ports on DS8884F storage systems, 64 ports on DS8884 storage systems, and a maximum of 128 ports on DS8886, DS8886F and DS8888F storage systems
- A maximum of 509 logins per Fibre Channel port
- A maximum of 8,192 logins per storage system
- A maximum of 1,280 logical paths on each Fibre Channel port
- Access to all 255 control-unit images (65,280 CKD devices) over each FICON port
- A maximum of 512 logical paths per control unit image

Note: IBM z13[®] and IBM z14[™] servers support 32,768 devices per FICON host channel, while IBM zEnterprise[®] EC12 and IBM zEnterprise BC12 servers support 24,576 devices per FICON host channel. Earlier IBM Z servers support 16,384 devices per FICON host channel. To fully access 65,280 devices, it is necessary to connect multiple FICON host channels to the storage system. You can access the devices through a Fibre Channel switch or FICON director to a single storage system FICON port.

The storage system supports the following operating systems for IBM Z 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).

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=s8g1S7000303>). 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 IBM Z 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 IBM Z 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 for SCSI end-to-end data protection on fixed block (FB) LUN volumes. This support applies to the IBM DS8880 unit (98x models). IBM Z support applies to FCP channels only.

IBM Z provides added end-to-end data protection between the operating system and the DS8880 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 DS8880 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 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 16 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, and 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 GiBs (2^{30})) or you can specify an ESS volume size (which is expressed in decimal GiBs (10^9) accurate to 0.1 GB). The unit of storage allocation for fixed block open systems volumes is one extent. The extent sizes for open volumes is either exactly 1 GiB, or 16 MiB. Any logical volume that is not an exact multiple of 1 GiB 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 TiB) in increments of one block. Supported sizes are from 1 to 16 TiB in increments of 1 GiB. 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 23 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 23. Capacity and models of disk volumes for IBM i hosts running IBM i operating system

Size	Protected model	Unprotected model
8.5 GB	A01	A81
17.5 GB	A02	A82
35.1 GB	A05	A85
70.5 GB	A04	A84
141.1 GB	A06	A86
282.2 GB	A07	A87
1 GB to 2000 GB	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 a CKD volume is either exactly a Mod-1 equivalent (which is 1113 cylinders), or it is 21 cylinders when using the small-extents option. 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, for a standard (non-EAV) 3390. The allocation of an EAV volume is expressed in increments of 3390 mod1 capacities (1113 cylinders) and can be expressed as integral multiples of 1113 between 65,667 - 1,182,006 cylinders or as the number of 3390 mod1 increments in the range of 59 - 1062.

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 z/OS. This support is available to you for the z/OS version 12.1, and later.

You can create a 1 TB IBM Z CKD volume. A IBM Z 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 IBM Z CKD volume, you must specify the number of cylinders that you want for the volume. The storage system and the z/OS have limits for the CKD EAV sizes. You can define CKD volumes with up to 1,182,006 cylinders, about 1 TB on the DS8880.

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

Quick initialization improves device initialization speed and allows a Copy Services relationship to be established after a device is created.

Quick volume initialization for IBM Z environments is supported. This support helps users who frequently delete volumes by reconfiguring capacity without waiting for initialization. Quick initialization initializes the data logical tracks or block 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.

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

Transparent cloud tiering

Transparent cloud tiering is a licensed function that enables volume data to be copied and transferred to cloud storage. DS8000 transparent cloud tiering is a feature in conjunction with z/OS and DFSMSHsm that provides server-less movement of archive and backup data directly to an object storage solution. Offloading the movement of the data from the host to the DS8000 unlocks DFSMSHsm efficiencies and saves z/OS CPU cycles.

DFSMSHsm has been the leading z/OS data archive solution for over 30 years. Its architecture is designed and optimized for tape, being the medium in which the data is transferred and archived.

Due to this architectural design point, there are inherent inefficiencies that consume host CPU cycles, including the following examples:

Movement of data through the host

All of the data must move from the disk through the host and out to the tape device.

Dual Data Movement

DSS must read the data from the disk and then pass the data from DSS to HSM, which then moves the data from the host to the tape.

16K block sizes

HSM separates the data within z/OS into small 16K blocks.

Recycle

When a tape is full, HSM must continually read the valid data from that tape volume and write it to a new tape.

HSM inventory

Reorgs, audits, and backups of the HSM inventory via the OCDS.

Transparent cloud tiering resolves these inefficiencies by moving the data directly from the DS8000 to the cloud object storage. This process eliminates the movement of data through the host, dual data movement, and the small 16K block size requirement. This process also eliminates recycle processing and the OCDS.

Transparent cloud tiering translates into significant savings in CPU utilization within z/OS, specifically when you are using both DFSMSHsm and transparent cloud tiering.

Modern enterprises adopted cloud storage to overcome the massive amount of data growth. The transparent cloud tiering system supports creating connections to cloud service providers to store data in private or public cloud storage. With transparent cloud tiering, administrators can move older data to cloud storage to free up capacity on the system. Point-in-time snapshots of data can be created on the system and then copied and stored on the cloud storage.

An external cloud service provider manages the cloud storage, which helps to reduce storage costs for the system. Before data can be copied to cloud storage, a connection to the cloud service provider must be created from the system. A cloud account is an object on the system that represents a connection to a cloud service provider by using a particular set of credentials. These credentials differ depending on the type of cloud service provider that is being specified. Most cloud service providers require the host name of the cloud service provider and an associated password, and some cloud service providers also require certificates to authenticate users of the cloud storage.

Public clouds use certificates that are signed by well-known certificate authorities. Private cloud service providers can use either self-signed certificate or a certificate that is signed by a trusted certificate authority. These credentials are defined on the cloud service provider and passed to the system through the administrators of the cloud service provider. A cloud account defines whether the system can successfully communicate and authenticate with the cloud service provider by using the account credentials. If the system is authenticated, it can then access cloud storage to either copy data to the cloud storage or restore data that is copied to cloud storage back to the system. The system supports one cloud account to a single cloud service provider. Migration between providers is not supported.

Client-side encryption for transparent cloud tiering ensures that data is encrypted before it is transferred to cloud storage. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system. You can use client-side encryption for transparent cloud tiering to download and decrypt data on any DS8000 storage system that uses the same set of key servers as the system that first encrypted the data.

Notes:

- Client-side encryption for transparent cloud tiering requires IBM Security Key Lifecycle Manager v3.0.0.2 or higher. For more information, see the IBM Security Key Lifecycle Manager online product documentation(www.ibm.com/support/knowledgecenter/SSWPVP/).
- Transparent cloud tiering supports the Key Management Interoperability Protocol (KMIP) only.

Cloud object storage is inherently multi-tenant, which allows multiple users to store data on the device, segregated from the other users. Each cloud service provider divides cloud storage into segments for each client that uses the cloud storage. These objects store only data specific to that client. Within the segment that is controlled by the user's name, DFSMSHsm and its inventory system controls the creation and segregation of containers that it uses to store the client data objects.

The storage system supports the OpenStack Swift and Amazon S3 APIs. The storage system also supports the IBM TS7700 as an object storage target and the following cloud service providers:

- Amazon S3

- IBM Bluemix - Cloud Object Storage
- OpenStack Swift Based Private Cloud

Dynamic volume expansion

Dynamic volume expansion is the capability 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 IBM Z 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: 16 TB
- IBM Z CKD volume types 3390 model 9 and custom: 65520 cylinders
- IBM Z CKD volume type 3390 model 3: 3339 cylinders
- IBM Z 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 path group. For FB volumes, *in use* means that the volumes are participating in a Copy Services relationship or there is 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.

Thin provisioning

Thin provisioning defines logical volume sizes that are larger than the physical capacity installed on the system. The volume allocates capacity on an as-needed basis as a result of host-write actions.

The thin provisioning feature enables the creation of extent space efficient logical volumes. Extent space efficient volumes are supported for FB and CKD volumes

and are supported for all Copy Services functionality, including FlashCopy targets where they provide a space efficient FlashCopy capability.

Releasing space on CKD volumes that use thin provisioning

On an IBM Z® host, the DFSMSdss SPACEREL utility can release space from thin provisioned CKD volumes that are used by either Global Copy or Global Mirror.

For Global Copy, space is released on the primary and secondary copies. If the secondary copy is the primary copy of another Global Copy relationship, space is also released on secondary copies of that relationship.

For Global Mirror, space is released on the primary copy after a new consistency group is formed. Space is released on the secondary copy after the next consistency group is formed and a FlashCopy commit is performed. If the secondary copy is the primary copy of another Global Mirror relationship, space is also released on secondary copies of that relationship.

Extent Space Efficient (ESE) capacity controls for thin provisioning

Use of thin provisioning can affect the amount of storage capacity that 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. Another thing that is 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.

An SNMP trap that is associated with the ESE capacity controls notifies you when the ESE extent usage in the pool exceeds an ESE extent threshold set by you. You are also notified 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.

IBM Easy Tier

Easy Tier is an optional feature that is provided at no cost. It can greatly increase the performance of your system by ensuring frequently accessed data is put on faster storage. 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 non-disruptively 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 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 DS8880.

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 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 feature 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 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 to help with the management of your ESE thin provisioning on fixed block (FB) or count key data (CKD) volumes.

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

Use the capabilities of Easy Tier to support:

Drive classes

The following drive classes are available, in order from highest to lowest performance. A pool can contain up to three drive classes.

Flash Tier 0 drives

The highest performance drives, which provide high I/O throughput and low latency.

Flash Tier 1 drives

The first tier of high capacity drives.

Flash Tier 2 drives

The second tier of high capacity drives.

Enterprise drives

SAS (10-K or 15-K RPM) disk drives.

Nearline drives

Nearline (7.2-K RPM) disk drives, which provide large data capacity but lower performance.

Three tiers

Using three tiers (each representing a separate drive class) and efficient algorithms improves system performance and cost effectiveness.

You can select from four drive classes to create up to three tiers. The drives within a tier must be homogeneous.

The following table lists the possible tier assignments for the drive classes. The tiers are listed according to the following values:

- | | |
|----------|---|
| 0 | Hot data tier, which contain the most active data. This tier can also serve as the home tier for new data allocations. |
| 1 | Mid-data tier, which can be combined with one or both of the other tiers and will contain data not moved to either of these tiers. This is by default the home tier for new data allocations. |
| 2 | Cold data tier, which contains the least active data. |

Table 24. Drive class combinations and tiers for systems with Flash Tier 0 drives as the highest performance drive class

Drive classes	Drive class combinations										
	Flash Tier 0	Flash Tier 0 + Flash Tier 1	Flash Tier 0 + Flash Tier 2	Flash Tier 0 + Flash Tier 1 + Flash Tier 2	Flash Tier 0 + Flash Tier 1 + Ent	Flash Tier 0 + Flash Tier 2 + Ent	Flash Tier 0 + Flash Tier 1 + NL	Flash Tier 0 + Flash Tier 2 + NL	Flash Tier 0 + Ent	Flash Tier 0 + Ent + NL	Flash Tier 0 + NL
Flash Tier 0	0	0	0	0	0	0	0	0	0	0	0
Flash Tier 1		1		1	1		1				
Flash Tier 2			1	2		1		1			
Enterprise (Ent)					2	2			1	1	
Nearline (NL)							2	2		2	1

Table 25. Drive class combinations and tiers for systems with Flash Tier 1 drives as the highest performance drive class

Drive classes	Drive class combinations						
	Flash Tier 1	Flash Tier 1 + Flash Tier 2	Flash Tier 1 + Flash Tier 2 + Ent	Flash Tier 1 + Flash Tier 2 + NL	Flash Tier 1 + Ent	Flash Tier 1 + Ent + NL	Flash Tier 1 + NL
Flash Tier 1	0	0	0	0	0	0	0
Flash Tier 2		1	1	1			
Enterprise (Ent)			2		1	1	
Nearline (NL)				2		2	1

Table 26. Drive class combinations and tiers for systems with Flash Tier 2 drives as the highest performance drive class

Drive classes	Drive class combinations			
	Flash Tier 2	Flash Tier 2 + Ent	Flash Tier 2 + Ent + NL	Flash Tier 2 + NL
Flash Tier 2	0	0	0	0
Enterprise (Ent)		1	1	
Nearline (NL)			2	1

Table 27. Drive class combinations and tiers for systems with Enterprise or Nearline drives as the highest performance drive class

Drive classes	Drive class combinations		
	Ent	Ent + NL	NL
Enterprise (Ent)	1	1	
Nearline (NL)		2	2

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 to the next lowest tier. 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.

Warm promotion

Active data that has higher IOPS is promoted to the next highest tier. Warm promotion is triggered whenever the lower tier is over its IOPS capacity. Selected warm extents are promoted to allow the lower tier to operate at its optimal load. Warm promotes 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 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 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 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 local system.

VMware vStorage API for Array Integration support

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

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

The following operations are supported:

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 storage system 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) 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 Copy Services 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 DS8880 with Fibre Channel adapters, see IBM DS8000 series online product documentation (http://www.ibm.com/support/knowledgecenter/ST5GLJ_8.1.0/com.ibm.storage.ssic.help.doc/f2c_securitybp.html) and select **Attaching and configuring hosts > 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, re protection, 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 (http://www.ibm.com/support/knowledgecenter/en/STAV45/hsg/hsg_vcplugin_kcwelcome_sonas.html) online documentation.

Performance for IBM Z

The storage system supports the following IBM performance enhancements for IBM Z environments.

- Parallel Access Volumes (PAVs)
- Multiple allegiance
- z/OS Distributed Data Backup
- z/HPF extended distance capability
- zHyperLink

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 IBM Z or S/390® hosts. PAV software support enables multiple users and jobs to simultaneously access a logical volume. Read and write operations can be accessed simultaneously to different domains. (The domain of an I/O operation is the specified extents to which the I/O operation applies.)

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) allows hosts, which are attached through a FICON interface, to access data on fixed block (FB) volumes through a device address on FICON interfaces.

If the zDDB LIC feature key is installed and enabled and a volume group type specifies either FICON 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 IBM Z host. The feature helps you manage multiple data protection environments and consolidate those into one environment that is managed by IBM Z. For more information, see “z/OS Distributed Data Backup” on page 122.

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 I/O ports. The storage system also supports the new zHPF I/O commands for multitrack I/O operations.

zHyperLink

zHyperLink is a short distance link technology that is designed for up to 10 times lower latency than zHPF. It can speed up transaction processing and improve active log throughput. zHyperLink is intended to complement FICON technology to accelerate I/O requests that are typically used for transaction processing.

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 or standard volume.
- PPRC supports any intermix of T10-protected or standard volumes. FlashCopy does not support intermix.
- PPRC supports copying from standard volumes to ESE volumes, or ESE volumes to Standard volumes, to allow migration with PPRC failover when both source and target volumes are on a DS8000 version 8.2 or higher.

The following Copy Services functions are available as optional features:

- Point-in-time copy, which includes IBM FlashCopy.
The FlashCopy function allows you to make point-in-time, full volume copies of data so that the copies are immediately available for read or write access. In IBM Z environments, you can also use the FlashCopy function to perform data set level copies of your data.
- Remote mirror and copy, which includes the following functions:
 - 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 IBM Z 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 IBM Z, as well as servers from Oracle and Hewlett-Packard.

You can manage these functions through a command-line interface that is called the DS CLI. 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 IBM Z environments, you can also use the FlashCopy function to perform data set level copies of your data. You can use the copy with standard backup tools that are available in your environment to create backup copies on tape.

FlashCopy is an optional function.

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 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 allows 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 71 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 if 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 storage system supports Incremental FlashCopy and Metro Global Mirror Incremental Resync on the same volume.

Safeguarded Copy

The Safeguarded Copy feature creates safeguarded backups that are not accessible by the host system and protects these backups from corruption that can occur in the production environment. You can define a Safeguarded Copy schedule to create multiple backups on a regular basis, such as hourly or daily. You can also restore a backup to the source volume or to a different volume. A backup contains the same metadata as the safeguarded source volume.

Safeguarded Copy can create backups with more frequency and capacity in comparison to FlashCopy volumes. The creation of safeguarded backups also impacts performance less than the multiple target volumes that are created by FlashCopy.

With backups that are outside of the production environment, you can use the backups to restore your environment back to a specified point in time. You can also extract and restore specific data from the backup or use the backup to diagnose production issues.

You cannot delete a safeguarded source volume before the safeguarded backups are deleted. The maximum size of a backup is 16 TB.

Copy Services Manager (available on the Hardware Management Console) is required to facilitate the use and management of Safeguarded Copy functions.

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.

Note: 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 location IBM System Storage Interoperation Center (SSIC) website (www.ibm.com/systems/support/storage/config/ssic).

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 allow 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 IBM Z 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 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 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 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 IBM Z and open systems data. Metro/Global Mirror uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

In a three-site Metro/Global Mirror, 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 needing 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 IBM Copy Services Manager.

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.

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 IBM Z environments. z/OS Global Mirror Multiple Reader helps maintain constant data consistency between mirrored sites and promotes efficient recovery. This function is supported on the storage system running in a IBM Z 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).

Global Copy is a non-synchronous long-distance copy option for data migration and backup.

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 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 on page 79 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 Copy Services features

The features of the Copy Services aid with planning for a disaster.

Table 28 provides a brief summary of the characteristics of the Copy Services features that are available for the storage system.

Table 28. Comparison of features

Feature	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 IBM Copy Services Manager
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 IBM Z host software	Nearly unlimited distance, highly scalable, and very low RPO.	Additional host server hardware and software is required. The RPO might grow if bandwidth capability is exceeded or host performance might be impacted.

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 through either the DS8000 Storage Management GUI or the DS CLI. If I/O Priority Manager is enabled on your storage system, you cannot use a zHyperLink connection.

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 29 lists performance groups that are predefined and have the associated performance policies:

Table 29. 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.		

Securing data

You can secure data with the encryption features that are supported by the storage system.

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 also requires access to at least two external key servers.

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 storage system automatically tests access to the encryption keys every 8 hours and access to the key servers every 5 minutes. You can verify access to key servers manually, initiate key retrieval, and monitor the status of attempts to access the key server.

Chapter 4. Planning the physical configuration

Physical configuration planning is your responsibility. Your technical support 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.

Administrative indicators

If applicable, models also include the following indicators:

- IBM / Openwave alliance
- IBM / EPIC attachment
- IBM systems, including System p and IBM Z
- Lenovo System x and BladeCenter
- 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
- Storage Appliance

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, including the licensed functions.

The Copy Services and z-Synergy Services licensed functions are based on usage requirements.

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 management console for each storage system. The management console feature code must be ordered for the base model within the storage system.
 - b. Decide whether a secondary management console is to be installed for the storage system. Adding a secondary 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:

- A primary management console
- A secondary management console

Primary and secondary 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 notifying 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 technical support representative to the HMC.

A secondary management console is available as an optional feature. The secondary 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 management console is included with every base frame along with a monitor and keyboard. An optional secondary management console is also available in the base frame.

Note: To preserve console function, the management consoles are not available as a general-purpose computing resource.

Feature codes for management consoles

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

Table 30. Feature codes for management consoles

Feature code	Description	Models
1141	Primary management console	A primary management consoles is required for model 984, 985, 986, and 988
1151	Secondary management console	Redundant management console for high availability For model 984, 985, 986, and 988, this feature is optional

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.

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
- Enclosure filler features
- Device adapter features

Storage enclosures and drives

DS8880 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.

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.

Note: Storage enclosures are installed in the frame from the bottom up.

Table 31 provide information on the placement of drive sets in the storage enclosure.

Table 31. 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

Feature codes for drive sets

Use these feature codes to order sets of encryption disk drives and flash drives for DS8880.

All drives that are installed in a standard drive enclosure pair or High Performance Flash Enclosure Gen2 pair must be of the same drive type, capacity, and speed.

The flash drives can be installed only in High Performance Flash Enclosures Gen2. See Table 34 on page 87 for the feature codes. Each High Performance Flash Enclosure Gen2 pair can contain 16, 32, or 48 flash drives. All flash drives in a High Performance Flash Enclosure Gen2 must be the same type and same capacity.

Table 32, Table 33 on page 87, Table 34 on page 87 list the feature codes for encryption drive sets based on drive size and speed.

Table 32. 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
5308	300 GB	2.5-in. disk drives	16	15 K	Yes	5 ¹ , 6, 10
5618	600 GB	2.5-in. disk drives	16	15 K	Yes	5 ¹ , 6, 10

Table 32. Feature codes for disk-drive sets (continued)

Feature code	Disk size	Drive type	Drives per set	Drive speed in RPM (K=1000)	Encryption drive	RAID support
5708	600 GB	2.5-in. disk drives	16	10 K	Yes	5 ¹ , 6, 10
5768	1.2 TB	2.5-in disk drives	16	10 K	Yes	6, 10
5778	1.8 TB	2.5-in disk drives	16	10 K	Yes	6, 10
5868	4 TB	3.5-in. NL disk drives	8	7.2 K	Yes	6, 10
5878	6 TB	3.5-in. NL disk drives	8	7.2 K	Yes	6, 10
Note: 1. RAID 5 is not supported for drives larger than 1 TB and requires a request for price quote (RPQ). For information, contact your service representative. 2. Drives are full disk encryption (FDE) self-encrypting drive (SED) capable.						

Table 33. Feature codes for flash-drive (SSD) sets for standard enclosures

Feature code	Disk size	Drive type	Drives per set	Drive speed in RPM (K=1000)	Encryption drive	RAID support
6158	400 GB	2.5-in flash drives	16	N/A	Yes	5 ¹ , 6, 10
6258	800 GB	2.5-in flash drives	16	N/A	Yes	5 ¹ , 6, 10
6358	1.6 TB	2.5-in flash drives	16	N/A	Yes	6, 10
Note: 1. RAID 5 is not supported for drives larger than 1 TB and requires a request for price quote (RPQ). For information, contact your service representative.						

Table 34. Feature codes for flash-drive sets for High Performance Flash Enclosures Gen2

Feature code	Disk size	Drive type	Drives per set	Drive speed in RPM (K=1000)	Encryption drive	RAID support
1610	400 GB	2.5-in. Flash Tier 0 drives	16	N/A	Yes	5, 6, 10
1611	800 GB	2.5-in. Flash Tier 0 drives	16	N/A	Yes	5, 6, 10
1612	1.6 TB	2.5-in. Flash Tier 0 drives	16	N/A	Yes	6, 10 ¹
1613	3.2 TB	2.5-in. Flash Tier 0 drives	16	N/A	Yes	6, 10 ¹
1623	3.8 TB	2.5-in. Flash Tier 1 drives	16	N/A	Yes	6, 10 ²
1624	7.6 TB	2.5-in. Flash Tier 2 drives	16	N/A	Yes	6 ³

Table 34. Feature codes for flash-drive sets for High Performance Flash Enclosures Gen2 (continued)

Feature code	Disk size	Drive type	Drives per set	Drive speed in RPM (K=1000)	Encryption drive	RAID support
Note: 1. RAID 5 is not supported for 1.6 TB and 3.2 TB Flash Tier 0 drives without a request for price quote (RPQ). For information, contact your sales representative. 2. RAID 5 is not supported for 3.8 TB Flash Tier 1 drives, and no RPQ is available. 3. RAID 5 and RAID 10 are not supported for 7.6 TB Flash Tier 2 drives, and no RPQ is available.						

Feature codes for storage enclosures

Use these feature codes to identify the type of drive enclosures for your storage system.

Table 35. Feature codes for storage enclosures

Feature code	Description	Models
1241	Standard drive-enclosure pair Note: This feature contains two filler sets in each enclosure.	984, 985, 986, 84E, 85E, 86E
1242	Standard drive-enclosure pair for 2.5-inch disk drives	984, 985, 986, 84E, 85E, 86E
1244	Standard drive-enclosure pair for 3.5-inch disk drives	984, 985, 986, 84E, 85E, 86E
1245	Standard drive-enclosure pair for 400 GB flash drives	984, 985, 986, 84E, 85E, 86E
1256	Standard drive-enclosure pair for 800 GB flash drives	984, 985, 986, 84E, 85E, 86E
1257	Standard drive-enclosure pair for 1.6 TB flash drives	984, 985, 986, 84E, 85E, 86E
1600	High Performance Flash Enclosure Gen2 pair with flash RAID controllers for 400 GB, 800 GB, 1.6 TB, 3.2 TB, 3.8 TB, and 7.6 TB flash drives	984, 985, 986, 988, 84E, 85E, 86E, 88E
1602	High Performance Flash Enclosure Gen2 pair for 400 GB, 800 GB, 1.6 TB, 3.2 TB, 3.8 TB, and 7.6 TB flash drives	988 and 88E
1604	Flash RAID adapter pair for 400 GB, 800 GB, 1.6 TB, 3.2 TB, 3.8 TB, and 7.6 TB flash drives	988 and 88E

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 Gen2, one filler feature provides a set of 16 fillers.

Feature codes for storage enclosure fillers

Use these feature codes to order filler sets for standard drive enclosures and High Performance Flash Enclosures Gen2.

Table 36. Feature codes for storage enclosures

Feature code	Description
2997	Filler set for 3.5-in. standard disk-drive enclosures; includes eight fillers
2999	Filler set for 2.5-in. standard disk-drive enclosures; includes 16 fillers
1699	Filler set for 2.5-in. High Performance Flash Enclosures Gen2; includes 16 fillers

Device adapters and flash RAID adapters

Device adapters and flash RAID adapters 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 RAID adapters 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.

Feature codes for device adapters

Use these feature codes to order device adapters for your storage system. Each feature includes two adapters. Device adapters are not supported for all-flash system type 533x.

Table 37. Feature codes for device adapters

Feature code	Description	Models
3053	4-port, 8 Gbps device adapter pair	984, 985, 986, 84E, 85E, 86E

Drive cables

If your storage system is a machine type 283x configuration, 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 A.
- If the disk drives connect to device adapters within the first expansion frame, order disk drive cable group B.
- If the disk drives are in a second expansion frame, order disk drive cable group C.
- If the disk drives are in a third expansion frame (85E and 86E only), order disk drive cable group D.
- If the disk drives are in a fourth expansion frame (85E and 86E only), order disk drive cable group E.

Feature codes for drive cables

If your storage system is a machine type 283x configuration, use these feature codes to order the cable groups.

If the disk drives are in a remote expansion frame (up to 20 meters from the base frame), order extended disk drive cable group C, D, or E as described here.

Table 38. Feature codes for drive cables

Feature code	Description	Connection Type
1246	Drive cable group A	(DS8886) Connects the drives to the device adapters within the same base model 985 or 986.
1247	Drive cable group B	(DS8886) Connects the drives to the device adapters in the first expansion model 85E or 86E.
1248	Drive cable group C	(DS8886) Connects the drives from a second expansion model 85E or 86E to the base model 985 or 986 and first expansion model 85E or 86E.
1249	Drive cable group D	(DS8886) Connects the drives from a third expansion model 85E or 86E to a second expansion model 85E or 86E.
1251	Drive cable group E	(DS8886) Connects the drives from a fourth expansion model 85E or 86E to a third expansion model 85E or 86E.
1252	Extended drive cable group C	(DS8886) Extended 20-meter cable that connects the drives from a second expansion model 85E or 86E to the base model 985, 986 and first expansion model 85E or 86E.
1253	Extended drive cable group D	(DS8886) Extended 20-meter cable that connects the drives from a third expansion model 85E or 86E to a second expansion model 85E or 86E.
1254	Extended drive cable group E	(DS8886) Extended 20-meter cable that connects the drives from a fourth expansion model 85E or 86E to a third expansion model 85E or 86E.
1261	Drive cable group A	(DS8884) Connects the disk drives to the device adapters within the same base model 984.

Table 38. Feature codes for drive cables (continued)

Feature code	Description	Connection Type
1262	Drive cable group B	(DS8884) Connects the disk drives to the device adapters in the first expansion model 84E.
1263	Drive cable group C	(DS8884) Connects the disk drives from a second expansion model 84E to the base model 984 and first expansion model 84E.
1266	Extended drive cable group C	(DS8884) Extended 20-meter cable that connects the drives from a second expansion model 84E to the base model 984 and first expansion model 84E.

Configuration rules for storage features

Use the following general configuration rules and ordering information to help you order storage features.

High Performance Flash Enclosures Gen2

Follow these configuration rules when you order storage features for storage systems with High Performance Flash Enclosures Gen2.

High Performance Flash Enclosures Gen2

For storage systems that support both standard drive enclosures and High Performance Flash Enclosures Gen2, DS8884 configurations can have one High Performance Flash Enclosure Gen2 pair in the base and first expansion frames, and DS8886 configurations can have two High Performance Flash Enclosure Gen2 pairs in the base and first expansion frames.

For all-flash storage systems, DS8884F configurations can have four High Performance Flash Enclosure Gen2 pair in the base frame, DS8886F configurations can have four High Performance Flash Enclosure Gen2 pair in the base and first expansion frame, and DS8888F configurations can have four High Performance Flash Enclosure Gen2 pair in the base frame, and six High Performance Flash Enclosure Gen2 pair in the first and second expansion frames.

Flash drive sets

Each High Performance Flash Enclosure Gen2 pair requires a minimum of one 16 flash-drive set.

Storage enclosure fillers

For High Performance Flash Enclosures Gen2, one filler feature provides a set of 16 fillers. If only one flash-drive set is ordered, then two storage enclosure fillers are needed to fill the remaining 32 slots in the High Performance Flash Enclosures Gen2 pair.

Flash RAID adapters

The flash RAID adapters are included in the I/O enclosure feature and are not ordered separately.

Drive cables

One drive cable set is required to connect the High Performance Flash Enclosure Gen2 to dedicated PCIe ports in the I/O enclosure. This cable set is included in the High Performance Flash Enclosure Gen2 feature and is not ordered separately.

Standard drive enclosures

For storage systems that support both standard drive enclosures and High Performance Flash Enclosures Gen2, follow these configuration rules when you order storage features for standard drive enclosures. All-flash systems do not support 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 populate 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 or 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.

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.

For DS8884, configurations with 64 GB of memory can have two device-adapter pairs, and configurations with 128 GB or 256 GB of memory can have four device-adapter pairs.

For DS8886, configurations with 8-core processors can have four device-adapter pairs, and configurations with 16-core and 24-core processors can have eight device-adapter pairs.

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 disk drives connect to device adapters within the same base frame, order disk drive cable group A.
- If the disk drives connect to device adapters within the first expansion frame, order disk drive cable group B.
- If the disk drives are in a second expansion frame, order disk drive cable group C.
- If the disk drives are in a third expansion frame (85E, 86E), order disk drive cable group D.
- If the disk drives are in a fourth expansion frame (85E, 86E), order disk drive cable group E.
- If the disk drives are in a remote expansion frame (up to 20 meters from the base frame), order extended disk drive cable group C, D, or E.

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 Gen2, there are 16 identical flash drives.

The logical configuration of your storage affects the effective capacity of the drive set.

Specifically, effective capacities vary depending on the following configurations:

RAID type and spares

Drives in the DS8000 must be configured as RAID 5, RAID 6, or RAID 10 arrays before they can be used, and then spare drives are assigned. 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.

Data format

Arrays are logically configured and formatted as fixed block (FB) or count key data (CKD) ranks. Data that is accessed by open systems hosts or Linux on IBM Z that support Fibre Channel protocol must be logically configured as FB. Data that is accessed by IBM Z hosts with z/OS or z/VM must be configured as CKD. Each RAID rank is divided into equal-sized segments that are known as extents.

The storage administrator has the choice to create extent pools of different extent sizes. The supported extent sizes for FB volumes are 1 GB or 16 MB and for CKD volumes it is one 3390 Mod1, which is 1113 cylinders or 21 cylinders. An extent pool cannot have a mix of different extent sizes.

On prior models of DS8000 series, a fixed area on each rank was assigned to be used for volume metadata, which reduced the amount of space available for use by volumes. In the DS8880 family, there is no fixed area for volume metadata, and this capacity is added to the space available for use. The metadata is allocated in the storage pool when volumes are created and is referred to as the pool overhead.

The amount of space that can be allocated by volumes is variable and depends on both the number of volumes and the logical capacity of these volumes. If thin provisioning is used, then the metadata is allocated for the entire volume when the volume is created, and not when extents are used, so over-provisioned environments have more metadata.

Metadata is allocated in units that are called metadata extents, which are 16 MB for FB data and 21 cylinders for CKD data. There are 64 metadata extents in each user extent for FB and 53 for CKD. The metadata space usage is as follows:

- Each volume takes one metadata extent.
- Ten extents (or part thereof) for the volume take one metadata extent.

For example, both a 3390-3 and a 3390-9 volume each take two metadata extents and a 128 GB FB volume takes 14 metadata extents.

Note: In a multiple tier pool volume, metadata is allocated on the upper tiers to provide maximum performance. Unless Nearline drives are the only tier in the pool, the metadata is never allocated on this tier. If Flash/SSD is in the pool, then metadata is allocated first on this tier and then on Enterprise drives (if available). A pool with 10% Flash/SSD or greater would have all of the volume metadata on this tier.

A simple way of estimating the maximum space that might be used by volume metadata is to use the following calculations:

FB Pool Overhead = (#volumes*2 + total volume extents / 10)/64 - rounded up to the nearest integer

CKD Pool Overhead = (#volumes*2 + total volume extents / 10)/53 - rounded up to the nearest integer

These calculations overestimate the space that is used by metadata by a small amount, but the precise details of each volume do not need to be known.

Examples:

- For an FB storage pool with 6,190 extents in which you expect to use thin provisioning and allocate up to 12,380 extents (2:1 overprovisioning) on 100 volumes, you would have a pool overhead of 23 extents -> $(100*2+12380/10)/64=22.46$.
- For a CKD storage pool with 6,190 extents in which you expect to allocate all the space on 700 volumes, then you would have a pool overhead of 39 extents -> $(700*2+6190/10)/53=38.09$.

RAID capacities for DS8880

Use the following information to calculate the physical and effective capacity for DS8880.

RAID capacities for High Performance Flash Enclosures Gen2

RAID 6 is the recommended and default RAID type for all drives over 1 TB. RAID 6 and RAID 10 are the only supported RAID types for 3.8 TB Flash Tier 1 drives. RAID 6 is the only supported RAID type for 7.6 TB Flash Tier 2 drives.

Table 39. RAID capacities for High Performance Flash Enclosures Gen2

Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives disk size	Physical capacity of Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives set	Rank type	Effective capacity of one rank in number of extents					
			RAID-10 arrays		RAID-5 arrays		RAID-6 arrays	
			3 + 3	4 + 4	6 + P	7 + P	5 + P + Q	6 + P + Q
400 GB	6.4 TB	FB Lg Ext	1049	1410	2132	2493	1771	2132
		FB Sm Ext	67170	90285	136507	159607	113393	136495
		CKD Lg Ext	1177	1582	2392	2797	1987	2392
		CKD Sm Ext	62388	83858	126797	148256	105328	126787
800 GB	12.8 TB	FB Lg Ext	2133	2855	4300	5023	3578	4300
		FB Sm Ext	136542	182781	275254	321475	229015	275239
		CKD Lg Ext	2392	3203	4823	5633	4013	4823
		CKD Sm Ext	126821	169768	255651	298601	212705	255655
1.6 TB	25.6 TB	FB Lg Ext	4301	5746	n/a	n/a	7191	8636
		FB Sm Ext	275284	367771	n/a	n/a	460243	552727
		CKD Lg Ext	4824	6445	n/a	n/a	8065	9686
		CKD Sm Ext	255684	341586	n/a	n/a	427475	513372

Table 39. RAID capacities for High Performance Flash Enclosures Gen2 (continued)

Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives disk size	Physical capacity of Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives set	Rank type	Effective capacity of one rank in number of extents					
			RAID-10 arrays		RAID-5 arrays		RAID-6 arrays	
			3 + 3	4 + 4	6 + P	7 + P	5 + P + Q	6 + P + Q
3.2 TB	51.2 TB	FB Lg Ext	8637	11527	n/a	n/a	14417	17307
		FB Sm Ext	552771	737753	n/a	n/a	922733	1107703
		CKD Lg Ext	9687	12928	n/a	n/a	16170	19412
		CKD Sm Ext	513414	685225	n/a	n/a	857029	1028843
3.8 TB	61.4 TB	FB Lg Ext	10371	13839	n/a	n/a	17308	20776
		FB Sm Ext	663766	885747	n/a	n/a	1107725	1329703
		CKD Lg Ext	11632	15522	n/a	n/a	19412	23302
		CKD Sm Ext	616506	822682	n/a	n/a	1028848	1235028
7.6 TB	123 TB	FB Lg Ext	n/a	n/a	n/a	n/a	34650	41587
		FB Sm Ext	n/a	n/a	n/a	n/a	2217663	2661631
		CKD Lg Ext	n/a	n/a	n/a	n/a	38863	46643
		CKD Sm Ext	n/a	n/a	n/a	n/a	2059760	2472118

RAID 5 array capacities for standard drive enclosures

The following table lists the RAID 5 array capacities for fully populated standard storage enclosures. RAID 5 is not supported for drives larger than 1 TB and requires a request for price quote (RPQ). For information, contact your sales representative.

Table 40. RAID capacities for RAID 5 arrays

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 5 array	
			6 + P	7 + P
300 GB disk drives	4,800	FB	1,664.30 (1,550)	1,947.77 (1,814)
		CKD	1,644.16 (1,738)	1,924.18 (2,034)
400 2.5-in. GB flash drives (SSD)	6,400	FB	2,289.22 (2,132)	2,676.84 (2,493)
		CKD	2,261.90 (2,391)	2,645.03 (2,796)
600 GB disk drives	9,600	FB	3,383.73 (3,156)	3,959.96 (3,688)
		CKD	3,348.86 (3,540)	3,912.68 (4,136)
800 GB 2.5-in. flash drive (SSD)	12,800	FB	4,617.09 (4,300)	5,392.33 (5,022)
		CKD	4,561.64 (4,822)	5,328.85 (5,633)

Table 40. RAID capacities for RAID 5 arrays (continued)

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 5 array	
			6 + P	7 + P
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 of DS8880 are different from rank capacities of earlier DS8000 storage systems. This change in the number of extents must be planned for when you move or migrate data to the DS8880.				

RAID 6 array capacities for standard drive enclosures

The following table lists the RAID 6 array capacities for fully populated storage enclosures.

Table 41. RAID capacities for RAID 6 arrays

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	1,347.55 (1,255)	1,628.87 (1,517)
		CKD	1,331.98 (1,408)	1,610.10 (1,702)
400 GB 2.5-in. flash drives (SSD)	6,400	FB	1,856.50 (1,729)	2,241.97 (2,088)
		CKD	1,834.30 (1,939)	2,214.60 (2,341)
600 GB disk drives	9,600	FB	2,750.93 (2,562)	3,318.94 (3,091)
		CKD	2,718.82 (2,874)	3,279.80 (3,467)
800 GB 2.5-in. flash drives (SSD)	12,800	FB	3,750.58 (3,493)	4,521.53 (4,211)
		CKD	3,706.45 (3,918)	4,467.98 (4,723)
1.2 TB disk drives	19,200	FB	5,541.58 (5,161)	6,676.53 (6,218)
		CKD	5,475.48 (5,788)	6,597.44 (6,974)
1.6 TB 2.5-in. flash drives (SSD)	25,600	FB	7,539.82 (7,022)	9,081.71 (8,458)
		CKD	7,450.74 (7,876)	8,974.75 (9,487)

Table 41. RAID capacities for RAID 6 arrays (continued)

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
1.8 TB disk drives	28,800	FB	8,331.16 (7,759)	10,034.12 (9,345)
		CKD	8,232.14 (8,702)	9,915.08 (10,481)
4 TB disk drives (Nearline)	64,000	FB	18,533.86 (17,261)	22,304.84 (20,773)
		CKD	18,313.72 (19,359)	220,40.98 (23,299)
6 TB disk drives (Nearline)	96,000	FB	27,896.89 (25,981)	33,573.76 (31,268)
		CKD	27,566.60 (29,140)	33,176.41 (35,070)
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 of DS8880 are different from rank capacities of earlier DS8000 storage systems. This change in the number of extents must be planned for when you move or migrate data to the DS8880.				

RAID 10 array capacities for standard drive enclosures

The following table lists the RAID 10 array capacities for fully populated storage enclosures.

Table 42. RAID capacities for RAID 10 arrays

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
300 GB disk drives	4,800	FB	813.90 (758)	1,097.36 (1,022)
		CKD	804.10 (850)	1,084.12 (1,146)
400 2.5-in. GB flash drives (SSD)	6,400	FB	1,126.36 (1,049)	1,513.98 (1,410)
		CKD	1,112.50 (1,176)	1,495.63 (1,581)
600 GB disk drives	9,600	FB	1,676.11 (1,561)	2,247.34 (2,093)
		CKD	1,656.46 (1,751)	2,220.27 (2,347)

Table 42. RAID capacities for RAID 10 arrays (continued)

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
800 GB 2.5-in. flash drive (SSD)	12,800	FB	2,290.29 (2,133)	3,065.53 (2,855)
		CKD	2,262.85 (2,392)	3,029.11 (3,202)
1.2 TB disk drives	19,200	FB	3,389.80 (3,157)	4,532.26 (4,221)
		CKD	3,349.81 (3,541)	4,478.39 (4,734)
1.6 TB flash drive (SSD)	25,600	FB	4,617.09 (4,300)	6,168.65 (5,745)
		CKD	4,562.58 (4,823)	6,096.06 (6,444)
1.8 TB disk drives	28,800	FB	5,103.49 (4,753)	6,817.19 (6,349)
		CKD	5,043.16 (5,331)	6,736.51 (7,121)
4 TB disk drives (Nearline)	64,000	FB	11,385.96 (10,604)	15,180.56 (14,138)
		CKD	11,250.84 (11,893)	15,000.81 (15,857)
6 TB disk drives (Nearline)	96,000	FB	17,136.92 (15,960)	22,849.23 (21,280)
		CKD	16,933.50 (17,900)	22,578.31 (23,867)
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 of DS8880 are different from rank capacities of earlier DS8000 storage systems. This change in the number of extents must be planned for when you move or migrate data to the DS8880.				

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
- Host adapters
- Host adapters Fibre Channel cables
- zHyperLink adapter
- zHyperLink cables
- Transparent cloud tiering adapters

- flash RAID adapters

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 or host adapters. Each I/O enclosure pair can support up to four device adapters (two pairs), and up to eight host adapters (not to exceed 32 host adapter ports).

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, up to four host adapters with eight ports, and up to eight host adapters with four ports. This feature includes four dedicated PCIe ports for attachment of High Performance Flash Enclosures Gen2.

Table 43. Feature codes for I/O enclosures

Feature code	Description
1303	I/O enclosure pair for PCIe group 3

Feature codes for I/O cables

Use these feature codes to order the I/O cables if your storage system is a DS8886. I/O cables are not required for other system types, including all-flash systems.

Table 44. 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.	985 and 986
1321	PCIe cable group 2	Connects device and host adapters in I/O enclosure pairs to the processor.	985 and 986

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 2105 (model 800 or 750) storage system. Fibre Channel host adapters are installed in an I/O enclosure.

Adapters are either 4-port or 8-port 8 Gbps, or 4-port 16 Gbps.

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 16 Gbps 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 8 Fibre Channel host adapters can be ordered with DS8884F. A maximum of 16 Fibre Channel host adapters can be ordered with DS8884 6-core and 12-core and DS8886 and DS8886F 8-core processors. A maximum of 32 Fibre Channel host adapters can be ordered with DS8886 and DS8886F 16-core and 24-core processors. A maximum of 16 Fibre Channel host adapters can be ordered with DS8888F with 24-core processors, and a maximum of 32 Fibre Channel host adapters can be ordered with DS8888F with 48-core processors.

Table 45. 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
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.		

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 46. 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)	Shortwave Fibre Channel or FICON host adapters (feature codes 3153, 3157, 3353)
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)	Longwave Fibre Channel or FICON adapters (feature codes 3253, 3257, 3453).
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 47. 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	50 m	35 m, but not recommended
OM3 (50 micron)	500 m	380 m	150 m	100 m
OM4 (50 micron)	500 m	400 m	190 m	125 m

zHyperLink adapters and cables

You can order zHyperLink adapters for your storage system configuration.

zHyperLink connections with IBM Z hosts provide low latency for random reads. Storage systems with 12-core and higher processors are supported. I/O Priority Manager must be disabled.

Note: The z-synergy Services license is required for zHyperLink.

Feature codes for zHyperLink I/O adapters

Use these feature codes to order zHyperLink I/O adapters for storage systems with 12-core and higher processors.

Each zHyperLink connection requires a zHyperLink I/O adapter to connect the zHyperLink cable to the storage system. Each zHyperLink I/O adapter card has one port, but you must order them in sets of two. Depending on the system configuration, zHyperLink I/O adapters can be ordered in quantities of 2, 4, 6, and 8.

Table 48. Feature codes for zHyperLink I/O adapters

Feature code	Description	Models
3500	1-port zHyperLink I/O adapter card	models 984 (requires feature code 4425), 985, 986, and 988 with 12-core and higher processors

Feature codes for zHyperLink cables

Use these feature codes to order cables to connect zHyperLink I/O adapters to storage systems with 12-core and higher processors . Take note of the distance capabilities for cable types.

Table 49. Feature codes for zHyperLink cables

Feature code	Cable type	Cable length	Compatible zHyperLink I/O adapter features
1450	OM4 50/125 micrometer, multimode, MTP connectors	40 m (131 ft)	zHyperLink I/O adapter (feature code 3500)
1451	OM4 50/125 micrometer, multimode, MTP connectors	150 m (492 ft)	

Feature codes for Transparent cloud tiering adapters

Use these feature codes to order adapter pairs to enhance Transparent cloud tiering connectivity for your storage system.

Transparent cloud tiering connectivity can be enhanced with 10 Gbps adapter pairs to improve bandwidth for a native cloud storage tier in IBM Z environments.

Table 50. Feature codes for Transparent cloud tiering adapter pairs

Feature code	Description	Models
3600	2-port 10 Gbps SFP+ optical/2-port 1 Gbps RJ-45 copper longwave adapter pair for 2U processor complex	model 984
3601	2-port 10 Gbps SFP+ optical/2-port 1 Gbps RJ-45 copper longwave adapter pair for 4U processor complex	models 985, 986, and 988

Feature codes for flash RAID adapters

Use these feature codes to order flash RAID adapters for a DS8888F storage system model 988 and 88E.

For your all-flash DS8888F storage system to have more than eight High Performance Flash Enclosure Gen2 pairs (feature code 1600), you must order a flash RAID adapter.

Table 51. Feature codes for flash RAID adapters

Feature code	Description	Models
1604	flash RAID adapter pair	all-flash model 988 and 88E

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, device adapters, and cards

Use these configuration rules and ordering information to help you order I/O enclosures, device adapters, and cards.

For storage systems that support both standard drive enclosures and High Performance Flash Enclosures Gen2, use the following tables to determine the number of I/O enclosures, device adapters, standard drive enclosures, and High Performance Flash Enclosure Gen2 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 enclosures, device adapters, standard drive enclosures, and High Performance Flash Enclosures Gen2 that you need in the storage system.

Table 52. Required I/O enclosures and device adapters for DS8884 (models 984 and 84E)

Processor type	Storage frame	Standard drive enclosure pair features (1241) ¹	Required device adapter pair features (3053) ²	High Performance Flash Enclosure Gen2 pair features (1600) ³	Required I/O enclosure features (1303) ⁴
6-core or 12-core ⁵	Base frame	0 - 4	1	0 - 1	1
6-core or 12-core ⁵	First expansion frame	1 - 4 5	1 1	0 0 - 1	0 1
6-core or 12-core ⁵	Second expansion frame	1 - 3 4 - 7	0 1	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 High Performance Flash Enclosure Gen2 feature represents one High Performance Flash Enclosure Gen2 pair, and a pair of flash RAID adapters and required cables. The maximum quantity is two High Performance Flash Enclosure Gen2 features for each I/O enclosure feature in the storage system.
4. 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 one High Performance Flash Enclosure Gen2 card pairs.
5. 12-core processors are supported only with zHyperLink.

Table 53. Required I/O enclosures and device adapters for DS8886 (models 985 and 85E)

Processor type	Storage frame	Standard drive enclosure pair features (1241) ¹	Required device adapter pair features (3053) ²	High Performance Flash Enclosure Gen2 pair features (1600) ³	Required I/O enclosure features (1303) ⁴
8-core, 16-core, or 24-core	Base frame	1	1	0 - 1	1
		2	2	0 - 1	1
		3	3	0 - 2	2
16-core or 24-core	First expansion frame	1	1	0 - 1	1
		2	2	0 - 1	1
		3	3	0 - 2	2
		4 - 5	4	0 - 2	2
16-core or 24-core	Second or third expansion frame	1	1	n/a	n/a
		2	2	n/a	n/a
		3	3	n/a	n/a
		4 - 9	4	n/a	n/a
16-core or 24-core	Fourth expansion frame	1	1	n/a	n/a
		2	2	n/a	n/a
		3	3	n/a	n/a
		4 - 6	4	n/a	n/a

Notes:

- Each storage enclosure feature represents one storage enclosure pair.
- 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.
- Each High Performance Flash Enclosure Gen2 feature represents one High Performance Flash Enclosure Gen2 pair, and a pair of flash RAID adapters and required cables. The maximum quantity is two High Performance Flash Enclosure Gen2 features for each I/O enclosure feature in the storage system.
- 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 one High Performance Flash Enclosure Gen2 card pairs.

Table 54. Required I/O enclosures and device adapters for DS8886 (models 986 and 86E)

Processor type	Storage frame	Standard drive enclosure pair features (1241) ¹	Required device adapter pair features (3053) ²	High Performance Flash Enclosure Gen2 pair features (1600) ³	Required I/O enclosure features (1303) ⁴
8-core, 16-core, or 24-core	Base frame	1	1	0 - 1	1
		2	2	0 - 1	1
16-core or 24-core	First expansion frame	1	1	0 - 1	1
		2	2	0 - 1	1
		3	3	0 - 2	2
		4	4	0 - 2	2

Table 54. Required I/O enclosures and device adapters for DS8886 (models 986 and 86E) (continued)

Processor type	Storage frame	Standard drive enclosure pair features (1241) ¹	Required device adapter pair features (3053) ²	High Performance Flash Enclosure Gen2 pair features (1600) ³	Required I/O enclosure features (1303) ⁴
16-core or 24-core	Second, third, or fourth expansion frame	1	1	n/a	n/a
		2	2	n/a	n/a
		3	3	n/a	n/a
		4 - 9	4	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 High Performance Flash Enclosure Gen2 feature represents one High Performance Flash Enclosure Gen2 pair, and a pair of flash RAID adapters and required cables. The maximum quantity is two High Performance Flash Enclosure Gen2 features for each I/O enclosure feature in the storage system.					
4. 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 one High Performance Flash Enclosure Gen2 card pairs.					

For all-flash storage systems, use the following tables to determine the number of I/O enclosures and High Performance Flash Enclosure Gen2 features that you need in various storage configurations. To use the table, find the rows that contain the type of all-flash storage system that you are configuring. Then, find the row that represents the number of High Performance Flash Enclosure Gen2 pairs that are installed in that storage system.

Table 55. Required I/O enclosures for all-flash DS8884F (model 984)

Processor type	Storage frame	High Performance Flash Enclosure Gen2 pair features (1600) ¹	Required I/O enclosure features (1303) ²
6-core or 12-core ³	Base frame	1 - 2 3 - 4	1 2
Notes:			
1. Each High Performance Flash Enclosure Gen2 feature represents one High Performance Flash Enclosure Gen2 pair, and a pair of flash RAID adapters and required cables. The maximum quantity is two High Performance Flash Enclosure Gen2 features for each I/O enclosure feature in the storage system.			
2. Each I/O enclosure feature represents an I/O enclosure pair. Each I/O enclosure pair can support two High Performance Flash Enclosure Gen2 pair.			
3. 12-core processors are supported only with zHyperLink.			

Table 56. Required I/O enclosures for all-flash DS8886F (models 985 and 85E, or 986 and 86E)

Processor type	Storage frame	High Performance Flash Enclosure Gen2 pair features (1600) ¹	Required I/O enclosure features (1303) ²
8-core, 16-core, or 24-core	Base frame	1 - 4	2
16-core, or 24-core	Expansion frame	1 - 4	2
Notes: - Each High Performance Flash Enclosure Gen2 feature represents one High Performance Flash Enclosure Gen2 pair, and a pair of flash RAID adapters and required cables. The maximum quantity is two High Performance Flash Enclosure Gen2 features for each I/O enclosure feature in the storage system. - Each I/O enclosure feature represents an I/O enclosure pair. Each I/O enclosure pair can support two High Performance Flash Enclosure Gen2 pair.			

Table 57. Required I/O enclosures for all-flash DS8888F (models 988 and 88E)

Processor type	Storage frame	High Performance Flash Enclosure Gen2 pair with flash RAID controllers features (1600) ¹	High Performance Flash Enclosure Gen2 pair features (1602) ³	Required I/O enclosure features (1303) ²
24-core or 48-core	Base frame	1 - 4	n/a	2
48-core	First expansion frame	1 - 4	1 - 2	2
48-core	Second expansion frame	n/a	1 - 6	n/a
Notes: - Each High Performance Flash Enclosure Gen2 feature represents one High Performance Flash Enclosure Gen2 pair, and a pair of flash RAID adapters and required cables. The maximum quantity is two High Performance Flash Enclosure Gen2 features for each I/O enclosure feature in the storage system. - Each I/O enclosure feature represents an I/O enclosure pair. Each I/O enclosure pair can support two High Performance Flash Enclosure Gen2 pair. - The first two feature 1602 in the first expansion frame each require a flash RAID adapter pair (feature 1604) in the base frame. The first two feature 1602 in the second expansion frame each require a flash RAID adapter pair (feature 1604) in the base frame. The last four feature 1602 in the second expansion frame each require a feature 1604 adapter in the first expansion frame.				

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

For storage systems that support both standard drive enclosures and High Performance Flash Enclosures Gen2, the following table lists the minimum and maximum host adapter features for the base frame (model 984, 985, 986, 988).

Table 58. 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
6-core or 12-core	Base frame + 2 expansion frames	2	16
8-core	Base frame	2	16
16-core or 24-core (model 985, 986)	Base frame + 1-4 expansion frames	2	32
Note: ¹ For all configurations, the maximum number of host adapters for any one frame cannot exceed 8 with 8-port 8-Gbps host adapters or 16 with 4-port 8-Gbps or 16-Gbps host adapters. You can add host adapters only to the base frame (model 984, 985, 986) and the first expansion frame (model 84E, 85E, 86E).			

For all-flash storage systems, the following table lists the minimum and maximum host adapter features for the base frame (model 984, 985, 986, 988).

Table 59. Minimum and maximum host adapter features for the all-flash base frame

All-flash 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
6-core or 12-core	Base frame	2	16
8-core	Base frame	2	16
16-core or 24-core (model 985, 986)	Base frame + 1 expansion frame	2	32
24-core (model 988)	Base frame	2	16
48-core (model 988)	Base frame + 2 expansion frame	2	32

Table 59. Minimum and maximum host adapter features for the all-flash base frame (continued)

All-flash 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
Note: ¹ For all configurations, the maximum number of host adapters for any one frame cannot exceed 8 with 8-port 8-Gbps host adapters or 16 with 4-port 8-Gbps or 16-Gbps host adapters. You can add host adapters only to the base frame (model 984, 985, 986, 988) and the first expansion frame (model 85E, 86E, 88E).			

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 a technical support representative moves existing host adapters from one slot to another, you can configure the host ports on your storage system by using the 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 98x) contain two processor enclosures (POWER8 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.

Table 60. Feature codes for processor licenses

Feature code	Description	Corequisite feature code for memory
4421	6-core POWER8 processor feature	4233, 4234, or 4235
4422	8-core POWER8 processor feature	4334 or 4335
4423	16-core POWER8 processor feature	4335 or 4336
4424	24-core POWER8 processor feature	4337 or 4338
4425	12-core POWER8 processor feature for zHyperLink support on model 984	4235
4878	24-core POWER8+ processor feature	4497
4898	48-core POWER8+ processor feature	4498

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 61. Feature codes for system memory

Feature code	Description	Model
4233 ¹	64 GB system memory	984 (6-core)
4234 ¹	128 GB system memory	984 (6-core)
4235 ²	256 GB system memory	984 (6-core and 12-core)
4334 ³	64 GB system memory	985 or 986 (8-core)
4335 ^{3 and 4}	128 GB system memory	985 or 986 (8-core and 16-core)
4336 ⁴	256 GB system memory	985 or 986 (16-core)
4337 ⁵	1 TB system memory	985 or 986 (24-core)
4338 ⁵	2 TB system memory	985 or 986 (24-core)
4497 ⁶	1 TB system memory	988 (24-core)
4498 ⁷	2 TB system memory	988 (48-core)

Table 61. Feature codes for system memory (continued)

Feature code	Description	Model
Notes: - Feature codes 4233 and 4234 require 6-core processor license feature code 4421. - Feature code 4235 requires either 6-core processor license feature code 4421 or 12-core processor license feature code 4425. 12-core processors are supported only with zHyperLink. - Feature codes 4334, and 4335 require 8-core processor license feature code 4422. - Feature codes 4335, and 4336 require 16-core processor license feature code 4423. - Feature codes 4337, and 4338 require 24-core processor license feature code 4424. - Feature code 4497 requires 24-core processor license feature code 4878. - Feature code 4498 requires 48-core processor license feature code 4898.		

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 DS8880 base frame and expansion frame, 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 DS8880 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: A minimum of one safety-approved ladder (feature code 1101) must be available at each installation site when the top exit bracket (feature code 1400) is specified for overhead cabling 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.

Attention: The regional designations are general. If the input voltage for the country uses a wye circuit, use the appropriate main power cables for EMEA (Europe, Middle East, and Africa) and Asia/Pacific. If input voltage for the country uses a delta circuit, use the appropriate main power cables for United States, Canada, and Latin America. For more information about electric currents for various countries, see the International Trade Administration website (<http://trade.gov/publications/abstracts/electric-current-abroad-2002.asp>).

Note: The IEC 60309 standard commercial/industrial pin and sleeve power connectors are often abbreviated "IEC '309" or simply "309 wall plug".

Table 62. Feature codes for power cords

Feature code	Power cord type	Wire gauge
1062	Single-phase power cord, 200-240V, 60A, 3-pin connector with IEC'309 service ready plug, for: United States, Canada, Latin America, Japan, and Taiwan (EMEA and other regions can optionally order) HBL360C6W, Pin and Sleeve Connector, IEC 60309, 2P3W HBL360R6W, AC Receptacle, IEC 60309, 2P3W	10 mm ² (7 AWG)
1063	Single-phase power cord, 200-240V, 63A, no connector (pre-stripped wires, termination ready for EMEA (except Japan), hard-wired, or uses client provided customer connector/plug as regionally allowed)	10 mm ² (7 AWG)
1086	Three-phase WYE (3ØY) voltage (five-wire 3+N+PE), 380-415V~ (nominal line-to-line (LL)), 30/32 A, IEC 60309 5-pin customer connector IEC'309 service ready plug, for EMEA (other regions under approved circumstances, such as United States, Canada, Latin America, Japan, and Taiwan can optionally order) HBL530C6V02, Pin and Sleeve Connector, IEC 60309, 4P5W HBL530R6V02, AC Receptacle, IEC 60309, 4P5W	6 mm ² (10 AWG)
1087	Three-phase Delta (3ØΔ) voltage (four-wire 3+PE (Protective Earth Ground)), 200-240V, 30 A, IEC 60309 4-pin customer connector with IEC'309 service ready plug, for: United States, Canada, Latin America, Japan, and Taiwan (EMEA and other regions can optionally order) HBL430C9W, Pin and Sleeve Connector, IEC 309, 3P4W HBL430P9V04, Pin and Sleeve Connector, IEC 309, 3P4W HBL430R9W, AC Receptacle, IEC 60309, 3P4W	6 mm ² (10 AWG)
1088	Three-phase WYE (3ØY) voltage (five-wire 3+N+PE), 380-415V~ (nominal line-to-line (LL)), 40 A, no customer connector provided (pre-stripped wires, termination ready for EMEA (except Japan), hard-wired, or uses client provided customer connector/plug as regionally allowed) Inline Connector: not applicable Receptacle: not applicable	6 mm ² (10 AWG)

Table 62. Feature codes for power cords (continued)

Feature code	Power cord type	Wire gauge
1089	Three-phase Delta (3ØΔ) voltage (four-wire 3+PE (Protective Earth Ground)), 200-240V, 60 A, IEC 60309 4-pin customer connector with IEC'309 service ready plug, for: United States, Canada, Latin America, Japan, and Taiwan (EMEA and other regions can optionally order) HBL460C9W, Pin and Sleeve Connector, IEC 309, 3P4W HBL460R9W, AC Receptacle, IEC 60309, 3P4W	10 mm ² (8 AWG)

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 includes a battery-service-module set.

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 40 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 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 these feature codes to order battery service modules for your base and expansion racks.

Table 63. Feature codes for the battery service module

Feature code	Description	Models
1050	Single-phase DC-UPS battery service module	984 and 84E, 985 and 85E
1052	Three-phase DC-UPS battery service module	986 and 86E, 988 and 88E

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 and expansion frame. The ePLD option protects the storage system from a power-line disturbance for 40 seconds.

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
- 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 40 seconds, rather than 4 seconds without the ePLD feature. This feature is optional for your storage system configuration.

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 64. Feature code for the ePLD

Feature code	Description	Models
1055	Extended power line disturbance	All

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 model 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 65. 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 technical 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 66. 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 533x or 904x machine function authorizations. However, the license functions are storage models. For example, the Base Function license is listed as a 533x or 904x model LF8. The 533x or 904x machine function authorization features are for billing purposes only.

The following licensed functions are available:

Base Function

The Base Function license is required for each storage system.

z-synergy Services

The z-synergy Services include z/OS licensed features that are supported on the storage system.

Copy Services

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

Copy Services Manager on Hardware Management Console

The Copy Services Manager on Hardware Management Console (CSM on HMC) license enables IBM Copy Services Manager to run on the Hardware Management Console, which eliminates the need to maintain a separate server for Copy Services functions.

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 licensed 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 283x or 904x 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

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. The Base Function, Copy Services, and Copy Services Manager on the Hardware Management Console licensed functions are available within this scope.

Count key data (CKD)

The function can be used only with data from FICON attached servers. The Copy Services, Copy Services Manager on the Hardware Management Console, and z-synergy Services licensed functions are available within this scope.

Both FB and CKD (ALL)

The function can be used with data from all attached servers. The Base Function, Copy Services, and Copy Services Manager on the Hardware Management Console licensed functions are available within this scope.

Some licensed functions have multiple license scope options, while other functions have only a single license scope.

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. A reduction of the license scope (changing ALL to FB or CKD) is a disruptive activity, which 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.

About this task

Licensed functions are purchased as function authorization features.

To order licensed functions, use the following general steps:

Procedure

1. Required. Order the Base Function license to support the total physical capacity of your storage system.

2. Optional. Order the z-synergy Services license to support the physical capacity of all ranks that are formatted as CKD.
3. Optional. Order the Copy Services license to support the total usable capacity of all volumes that are involved in one or more copy services functions.

Note: The Copy Services license is based on the usable capacity of volumes and not on physical capacity. If overprovisioning is used on the DS8880 with a significant amount of Copy Services functionality, then the Copy Services license needs only to be equal to the total rank usable capacity (even if the logical volume capacity of volumes in Copy Services is greater). For example, if the total rank usable capacity of a DS8880 is 100 TB but there are 200 TB of thin provisioning volumes in Metro Mirror, then only a 100 TB of Copy Services license is needed.

4. Optional. Order the Copy Services Manager on the Hardware Management Console license that support the total usable capacity of all volumes that are involved in one or more copy services functions.

Rules for ordering licensed functions

A Base Function license is required for every base frame. All other licensed functions are optional and must have a capacity that is equal to or less than the Base Function license.

For all licensed functions, you can combine feature codes to order the exact capacity that you need. For example, if you require 160 TB of Base Function license capacity, order 10 of feature code 8151 (10 TB each up to 100 TB capacity) and 4 of feature code 8152 (15 TB each, for an extra 60 TB).

When you calculate usable capacity for the Copy Services license, use the size of each volume involved in a copy services relationship and multiply by the size of each extent.

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 table to determine the total size of each regular drive feature in your storage system, and then add all the values.

Table 67. Total physical capacity for drive-set features

Drive sizes	Total physical capacity	Drives per feature
300 GB disk drives	4.8 TB	16
400 GB flash drives	6.4 TB	16
400 GB flash drives	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
1.2 TB disk drives	19.2 TB	16
1.6 TB flash drives	25.6 TB	16
1.8 TB disk drives	28.8 TB	16
3.2 TB flash drives	51.2 TB	16
3.8 TB flash drives	61.4	16
4 TB disk drives	32 TB	8
6 TB disk drives	48 TB	8

Table 67. Total physical capacity for drive-set features (continued)

Drive sizes	Total physical capacity	Drives per feature
7.6 TB flash drives	122 TB	16

Rules for removing a licensed function

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 is a nondisruptive activity but takes effect at the next machine IML.

If you have a licensed function and no longer want to use it, you can deactivate the license 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 nondisruptive 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).

Base Function license

The Base Function license provides essential functions for your storage system. A Base Function license is required for each storage system.

The Base Function license is available for the following license scopes: FB and ALL (both FB and CKD).

The Base Function license includes the following features:

- Database Protection
- Encryption Authorization
- Easy Tier
- I/O Priority Manager
- Operating Environment License (OEL)
- Thin Provisioning

The Base Function license feature codes are ordered in increments up to a specific capacity. For example, if you require 160 TB of capacity, order 10 of feature code 8151 (10 TB each up to 100 TB capacity) and 4 of feature code 8152 (15 TB each, for an extra 60 TB).

The Base Function license includes the following feature codes.

Table 68. Base Function license feature codes

Feature Code	Feature code for licensed function indicator
8151	BF - 10 TB (up to 100 TB capacity)
8152	BF - 15 TB (from 100.1 TB to 250 TB capacity)
8153	BF - 25 TB (from 250.1 TB to 500 TB capacity)
8154	BF - 75 TB (from 500.1 to 1250 TB capacity)
8155	BF - 175 TB (from 1250.1 TB to 3000 TB capacity)
8156	BF - 300 TB (from 3000.1 TB to 6000 TB capacity)
8160	BF - 500 TB (from 6000.1 TB to 10,000 TB capacity)

Base Function license rules

The Base Function license authorizes you to use the model configuration at a specific capacity level. The Base Function license must cover the full physical capacity of your storage system, which includes the physical capacity of any expansion frames within the storage system. The license capacity must cover both open systems data (fixed block data) and IBM Z data (count key data). All other licensed functions must have a capacity that is equal to or less than the Base Function license.

Note: Your storage system cannot be logically configured until you activate the Base Function license. On activation, drives can be logically configured up to the extent of the Base Function license authorization level.

As you add more drives to your storage system, you must increase the Base Function license authorization level for the storage system by purchasing more license features. Otherwise, you cannot logically configure the additional drives for use.

Database Protection

The IBM Database Protection feature provides the highest level of protection for Oracle databases by detecting corrupted Oracle data and preventing it from being processed to storage.

The IBM Database Protection feature complies with the Oracle Hardware Assisted Resilient Data (HARD) initiative, which provides an end-to-end data protection between an Oracle database and permanent storage devices.

Data must pass through many software and hardware layers on its way to storage. It is possible for the data to become corrupted, on a rare occasion, caused by a malfunction in an intermediate layer. With the IBM Database Protection feature, an IBM DS8000 model can validate whether Oracle data blocks are consistent using the same logic that Oracle uses. This validation is done before the write request is processed. You can designate how the transaction is managed: either rejected and reported, or processed and reported.

Encryption Authorization

The Encryption Authorization feature provides data encryption by using IBM Full Disk Encryption (FDE) and key managers, such as IBM Security Key Lifecycle Manager.

The Encryption Authorization feature secures data at rest and offers a simple, cost-effective solution for securely erasing any disk drive that is being retired or re-purposed (cryptographic erasure). The storage system uses disks that have FDE encryption hardware and can perform symmetric encryption and decryption of data at full disk speed with no impact on performance.

IBM Easy Tier

Support for IBM Easy Tier is available with the IBM Easy Tier feature.

The Easy Tier feature enables the following modes:

- Easy Tier: automatic mode
- Easy Tier: manual mode

The 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

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.

Operating environment license

The operating environment model and features establish the extent of IBM authorization for the use of the IBM DS operating environment.

To determine the operating environment license support function, see “Machine types overview” on page 4.

Thin provisioning

Thin provisioning defines logical volume sizes that are larger than the physical capacity installed on the system. The volume allocates capacity on an as-needed basis as a result of host-write actions.

The thin provisioning feature enables the creation of extent space efficient logical volumes. Extent space efficient volumes are supported for FB and CKD volumes and are supported for all Copy Services functionality, including FlashCopy targets where they provide a space efficient FlashCopy capability.

z-synergy Services license

The z-synergy Services license includes z/OS® features that are supported on the storage system.

The z-synergy Services license is available for the following license scope: CKD.

The z-synergy Services license includes the following features:

- High Performance FICON for z Systems
- HyperPAV
- Parallel Access Volumes (PAV)
- Transparent cloud tiering
- z/OS Distributed Data Backup
- zHyperLink

The z-synergy Services license also includes the ability to attach FICON channels.

The z-synergy Services license feature codes are ordered in increments up to a specific capacity. For example, if you require 160 TB of capacity, order 10 of feature code 8351 (10 TB each up to 100 TB capacity), and 4 of feature code 8352 (15 TB each, for an extra 60 TB).

The z-synergy Services license includes the feature codes listed in the following table.

Table 69. z-synergy Services license feature codes

Feature Code	Feature code for licensed function indicator
8350	zsS - inactive
8351	zsS - 10 TB (up to 100 TB capacity)
8352	zsS - 15 TB (from 100.1 TB to 250 TB capacity)
8353	zsS - 25 TB (from 250.1 TB to 500 TB capacity)
8354	zsS - 75 TB (from 500.1 to 1250 TB capacity)
8355	zsS - 175 TB (from 1250.1 TB to 3000 TB capacity)
8356	zsS - 300 TB (from 3000.1 TB to 6000 TB capacity)
8360	zsS - 500 TB (from 6000.1 TB to 10,000 TB capacity)

z-synergy Services license rules

A z-synergy Services license is required for only the total physical capacity that is logically configured as count key data (CKD) ranks for use with IBM Z host systems.

Note: If z/OS Distributed Data Backup is being used on a system with no CKD ranks, a 10 TB z-synergy Services license must be ordered to enable the FICON attachment functionality.

High Performance FICON for z Systems

High Performance FICON for z Systems (zHPF) is an enhancement to the IBM FICON architecture to offload I/O management processing from the z Systems channel subsystem to the DS8880 Host Adapter and controller.

zHPF is an optional feature of z Systems server and of the DS8880. Recent enhancements to zHPF include Extended Distance Facility zHPF List Pre-fetch support for IBM DB2® and utility operations, and zHPF support for sequential access methods. All of DB2 I/O is now zHPF-capable.

IBM HyperPAV

IBM HyperPAV associates the volumes with either an alias address or a specified base logical volume number. When a host system requests IBM HyperPAV processing and the processing is enabled, aliases on the logical subsystem are placed in an IBM HyperPAV alias access state on all logical paths with a given path group ID.

Parallel Access Volumes

The parallel access volumes (PAV) features establish the extent of IBM authorization for the use of the parallel access volumes function.

Parallel Access Volumes (PAVs), also referred to as aliases, provide your system with access to volumes in parallel when you use an IBM Z host.

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

Transparent cloud tiering

Transparent cloud tiering provides a native cloud storage tier for IBM Z environments. Transparent cloud tiering moves data directly from the storage system to cloud object storage, without sending data through the host.

Transparent cloud tiering provides cloud object storage (public, private, or on-premises) as a secure, reliable, transparent storage tier that is natively integrated with the storage system. Transparent cloud tiering on the storage system is fully integrated with DFSMSHsm, which reduces CPU utilization on the host when you are migrating and recalling data in cloud storage. You can use the IBM Z host to manage transparent cloud tiering and attach metadata to cloud objects.

The storage system supports the OpenStack Swift and Amazon S3 APIs. The storage system also supports the IBM TS7700 as an object storage target and the following cloud service providers:

- Amazon S3
- IBM Bluemix - Cloud Object Storage
- OpenStack Swift Based Private Cloud

z/OS Distributed Data Backup

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

If zDDB is installed and enabled and a volume group type specifies either FICON 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.

zHyperLink

zHyperLink is a short distance link technology that is designed for up to 10 times lower latency than zHPF.

zHyperLink can speed up transaction processing and improve active log throughput.

zHyperLink is intended to complement FICON technology to accelerate I/O requests that are typically used for transaction processing.

Copy Services license

Copy Services features help you implement storage solutions to keep your business running 24 hours a day, 7 days a week by providing data duplication, data migration, and disaster recovery functions. The Copy Services license is based on usable capacity of the volumes involved in Copy Services functionality.

The Copy Services license is available for the following license scopes: FB and ALL (both FB and CKD).

The Copy Services license includes the following features:

- Global Mirror
- Metro Mirror
- Metro/Global Mirror
- Point-in-Time Copy/FlashCopy
- Safeguarded Copy
- z/OS Global Mirror
- z/OS Metro/Global Mirror Incremental Resync (RMZ)

The Copy Services license feature codes are ordered in increments up to a specific capacity. For example, if you require 160 TB of capacity, order 10 of feature code 8251 (10 TB each up to 100 TB capacity), and 4 of feature code 8252 (15 TB each, for an extra 60 TB).

The Copy Services license includes the following feature codes.

Table 70. Copy Services license feature codes

Feature Code	Feature code for licensed function indicator
8250	CS - inactive
8251	CS - 10 TB (up to 100 TB capacity)
8252	CS - 15 TB (from 100.1 TB to 250 TB capacity)
8253	CS - 25 TB (from 250.1 TB to 500 TB capacity)
8254	CS - 75 TB (from 500.1 to 1250 TB capacity)
8255	CS - 175 TB (from 1250.1 TB to 3000 TB capacity)
8256	CS - 300 TB (from 3000.1 TB to 6000 TB capacity)
8260	CS - 500 TB (from 6000.1 TB to 10,000 TB capacity)

Copy Services license rules

The following ordering rules apply when you order the Copy Services license:

- The Copy Services license should be ordered based on the total usable capacity of all volumes involved in one or more Copy Services relationships.
- The licensed authorization must be equal to or less than the total usable capacity allocated to the volumes that participate in Copy Services operations.
- You must purchase features for both the source (primary) and target (secondary) storage system.

Remote mirror and copy functions

The Copy Services license establishes the extent of IBM authorization for the use of the remote mirror and copy functions on your storage system.

The following functions are included:

- Metro Mirror
- Global Mirror
- Global Copy
- Metro/Global Mirror
- Multiple Target PPRC

FlashCopy function (point-in-time copy)

FlashCopy creates a copy of a source volume on the target volume. This copy is called a point-in-time copy.

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

Safeguarded Copy

The Safeguarded Copy feature, available with the Copy Services license, creates backups of data that you can restore to the source volume or a different volume.

The Safeguarded Copy feature creates safeguarded backups that are not accessible by the host system and protects these backups from corruption that can occur in the production environment. You can define a Safeguarded Copy schedule to create multiple backups on a regular basis, such as hourly or daily. You can also restore a backup to the source volume or to a different volume. A backup contains the same metadata as the safeguarded source volume.

Safeguarded Copy can create backups with more frequency and capacity in comparison to FlashCopy volumes. The creation of safeguarded backups also impacts performance less than the multiple target volumes that are created by FlashCopy.

With backups that are outside of the production environment, you can use the backups to restore your environment back to a specified point in time. You can also extract and restore specific data from the backup or use the backup to diagnose production issues.

You cannot delete a safeguarded source volume before the safeguarded backups are deleted. The maximum size of a backup is 16 TB.

z/OS Global Mirror

z/OS Global Mirror (previously known as Extended Remote Copy or XRC) provides a long-distance remote copy solution across two sites for open systems and IBM Z data with asynchronous technology.

z/OS Metro/Global Mirror Incremental Resync

z/OS Metro/Global Mirror Incremental Resync (RMZ) is an enhancement for z/OS 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 Global Mirror configurations.

The storage system supports z/OS 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.

Copy Services Manager on the Hardware Management Console license

IBM Copy Services Manager 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). IBM Copy Services Manager is available on the Hardware Management Console (HMC), which eliminates the need to maintain a separate server for Copy Services functions.

The Copy Services Manager on Hardware Management Console (CSM on HMC) license is available for the following license scopes: FB and ALL (both FB and CKD).

The CSM on HMC license feature codes are ordered in increments up to a specific capacity. For example, if you require 160 TB of capacity, order 10 of feature code 8451 (10 TB each up to 100 TB capacity), and 4 of feature code 8452 (15 TB each, for an extra 60 TB).

The CSM on HMC license includes the following feature codes.

Table 71. Copy Services Manager on Hardware Management Console license feature codes

Feature Code	Feature code for licensed function indicator
8450	CSM on HMC - inactive
8451	CSM on HMC - 10 TB (up to 100 TB capacity)
8452	CSM on HMC - 15 TB (from 100.1 TB to 250 TB capacity)
8453	CSM on HMC - 25 TB (from 250.1 TB to 500 TB capacity)
8454	CSM on HMC - 75 TB (from 500.1 to 1250 TB capacity)
8455	CSM on HMC - 175 TB (from 1250.1 TB to 3000 TB capacity)
8456	CSM on HMC - 300 TB (from 3000.1 TB to 6000 TB capacity)
8460	CSM on HMC - 500 TB (from 6000.1 TB to 10,000 TB capacity)

Chapter 6. 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 equipment and documents,” on page 197.

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 a technical 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 technical 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 (racks and frames) must be gradually acclimated to the surrounding environment to prevent condensation.

When server and storage equipment (racks and frames) is shipped in a climate where the outside temperature is below the dew point of the destination (indoor location), there is a possibility that water condensation can form on the cooler inside and outside surfaces of the equipment when the equipment is brought indoors.

Sufficient time must be allowed for the shipped equipment to gradually reach thermal equilibrium with the indoor environment before you remove the shipping bag and energize the equipment. Follow these guidelines to properly acclimate your equipment:

- Leave the system in the shipping bag. If the installation or staging environment allows it, leave the product in the full package to minimize condensation on or within the equipment.

- Allow the packaged product to acclimate for 24 hours.¹ If there are visible signs of condensation (either external or internal to the product) after 24 hours, acclimate the system without the shipping bag for an additional 12 - 24 hours or until no visible condensation remains.
- Acclimate the product away from perforated tiles or other direct sources of forced air convection to minimize excessive condensation on or within the equipment.

¹ Unless otherwise stated by product-specific installation instructions.

Note: Condensation is a normal occurrence, especially when you ship equipment in cold-weather climates. All IBM® products are tested and verified to withstand condensation that is produced under these circumstances. When sufficient time is provided to allow the hardware to gradually acclimate to the indoor environment, there should be no issues with long-term reliability of the product.

Shipment weights and dimensions

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 has a plywood front door on the package, with a corrugated paperboard outer wrap. 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.

Table 72 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.

Table 72. Packaged dimensions and weight for storage system frames (all countries)

Container	Packaged dimensions	Maximum packaged weight
DS8880 base frame (model 984)	Height 2.08 m (81.9 in.) Width 0.95 m (37.4 in.) Depth 1.50 m (59.1 in.)	1065 kg (2348 lb)
DS8880 base frame (model 985)	Height 2.08 m (81.9 in.) Width 0.95 m (37.4 in.) Depth 1.50 m (59.1 in.)	1250.5 kg (2757 lb)
DS8880 base frame (model 986)	Height 2.08 m (81.9 in.) Width 0.95 m (37.4 in.) Depth 1.50 m (59.1 in.)	1187.5 kg (2618 lb)

Table 72. Packaged dimensions and weight for storage system frames (all countries) (continued)

Container	Packaged dimensions	Maximum packaged weight
DS8880 base frame (model 988)	Height 2.08 m (81.9 in.) Width 0.95 m (37.4 in.) Depth 1.50 m (59.1 in.)	1184 kg (2610 lb)
DS8880 expansion frame (model 84E) ¹	Height 2.08 m (81.9 in.) Width 0.95 m (37.4 in.) Depth 1.50 m (59.1 in.)	1128 kg (2487 lb)
DS8880 expansion frame (model 85E)	Height 2.08 m (81.9 in.) Width 0.95 m (37.4 in.) Depth 1.50 m (59.1 in.)	1205 kg (2657 lb)
DS8880 expansion frame (model 86E)	Height 2.08 m (81.9 in.) Width 0.95 m (37.4 in.) Depth 1.50 m (59.1 in.)	1043 kg (2297 lb)
DS8880 expansion frame (model 88E)	Height 2.08 m (81.9 in.) Width 0.95 m (37.4 in.) Depth 1.50 m (59.1 in.)	1070 kg (2359 lb)
Note: 1. The all-flash DS8884F storage system configuration does not support an expansion frame. 2. The top expansion (feature code 0170) is shipped separately.		


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.

Before you begin

Ensure you read the following caution when you position the rack. If you are relocating the frame, ease it out of its current position and pull out the outriggers for the remaining major part of the relocation. Roll the rack on its castors until you get close to its intended location. Keep the supplemental outriggers in position as shown in the following image. When the rack is near the final location, you can recede the outriggers back into the recessed position, flush with the outsides of the rack. The outriggers are only intended to help move the rack and are not intended to support the rack in its final location. To prevent unintended movement and ensure stability of the rack, you can put down the leveler jacks.

CAUTION: The rack cabinet is supplied with native built-in extendable outriggers with small floating supplemental castors as motion anti-tip features. They must all be extended into a latched position before and during cabinet movement or relocation. These native built-in outriggers must not be removed completely, but rather recessed in when finished to ensure they are readily available for future use. (C050)



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?

Procedure

Use the following steps to ensure that your planned installation site meets space and floor load requirements:

1. Identify the base frame and expansion frames that are included in your storage system.
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.). 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 2844 kg (6270 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 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 plyon) 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 a 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 134 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 and internal cable routing brackets for rack 1 or rack 2. The following notes provide more information about the color-coded cable routing and components in Figure 14 on page 134.

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 the 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 a partner storage system are also located here.

3 Mainline power cords. Two top-exit mainline power cords for each rack, which are shown in green, are routed here.

Notes:

- A technical service representative 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 a technical service representative.
- All remaining cables are internally routed and connected by a technical 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 73. 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 technical 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 5-foot ladder. This ladder provides technical service representatives the ability to perform power safety checks and other service activities on the top of your storage system. Without this approved ladder, technical 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 enclosures.

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 DS8880 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 DS8880.




Figure 15. Cable cutouts for DS8880

2. If you install the storage system on a raised floor, use the following measurements when you cut the floor tile for cabling:
 - **1** Width: 41.91 cm (16.5 in.)
 - **2** End of frame to edge of cable cutout: 10.0 cm (3.9 in)
 - **3** Depth: 8.89 cm (3.5 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

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** and **2** Length of tailgate opening
- **3** Width of tailgate opening




Figure 16. DS8880 top exit bracket feature

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 17 on page 138

- 1** Front service clearance
- 2** Depth of frame without covers
- 3** Width of frame without covers
- 4** Minimum dimension between casters and outside edges of frames
- 5** Leveling pads
- 6** Back service clearance
- 7** Leveling pads




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).
 - Use the table to determine the required side, front, and rear weight distribution area for the specified floor load rating. The side dimensions for the weight distribution area have a maximum of 76.2 cm (30 in). If the side dimensions required to meet a specified floor load rating are greater than the maximum allowed, then the floor load rating is not listed in the table.
 - The maximum per caster weight that is transferred 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: If you are unsure about the correct placement and weight distribution areas for your storage system consult a structural engineer.

Table 74. Floor load ratings and required weight-distribution areas

Configuration _a	Total weight of configuration _b	Floor load rating, kg per m ² (lb per ft ²)	Weight distribution areas _{c, d, e}		
			Sides cm (in.)	Front cm (in.)	Rear cm (in.)
Model 984	976 kg (2151 lb)	610 (125)	7.62 (3)	76.2 (30)	76.2 (30)
		488 (100)	18 (7)	76.2 (30)	76.2 (30)
		439 (90)	25 (10)	76.2 (30)	76.2 (30)
		342 (70)	51 (20)	76.2 (30)	76.2 (30)
Model 984 and one 84E expansion model ⁶	2014 kg (4441 lb)	610 (125)	15.3 (6)	76.2 (30)	76.2 (30)
		488 (100)	38 (15)	76.2 (30)	76.2 (30)
		439 (90)	56 (22)	76.2 (30)	76.2 (30)
		342 (70)	not supported	not supported	not supported
Model 984 and two 84E expansion models ⁶	2863 kg (6311 lb)	610 (125)	15.25 (6)	76.2 (30)	76.2 (30)
		488 (100)	51 (20)	76.2 (30)	76.2 (30)
		439 (90)	74 (29)	76.2 (30)	76.2 (30)
		342 (70)	not supported	not supported	not supported

Table 74. Floor load ratings and required weight-distribution areas (continued)

Configuration _a	Total weight of configuration _b	Floor load rating, kg per m ² (lb per ft ²)	Weight distribution areas _{c, d, e}		
			Sides cm (in.)	Front cm (in.)	Rear cm (in.)
Model 985	1144 kg (2520 lb)	610 (125)	13 (5)	76.2 (30)	76.2 (30)
		488 (100)	28 (11)	76.2 (30)	76.2 (30)
		439 (90)	38 (15)	76.2 (30)	76.2 (30)
		342 (70)	66 (26)	76.2 (30)	76.2 (30)
Model 985 and first 85E expansion model	2259 kg (4980 lb)	610 (125)	23 (9)	76.2 (30)	76.2 (30)
		488 (100)	53.4 (21)	76.2 (30)	76.2 (30)
		439 (90)	69 (27)	76.2 (30)	76.2 (30)
		342 (70)	not supported	not supported	not supported
Model 985 with first and second 85E expansion models (with top expansion) ⁶	3190 kg (7032 lb)	610 (125)	28 (11)	76.2 (30)	76.2 (30)
		488 (100)	69 (27)	76.2 (30)	76.2 (30)
		439 (90)	not supported	not supported	not supported
		342 (70)	not supported	not supported	not supported
Model 985 with first, second, and third 85E expansion models (with top expansion) ⁶	4121 kg (9084 lb)	610 (125)	33 (13)	76.2 (30)	76.2 (30)
		488 (100)	not supported	not supported	not supported
		439 (90)	not supported	not supported	not supported
		342 (70)	not supported	not supported	not supported
Model 985 with first, second, third, and fourth 85E expansion models (with top expansion) ⁶	4933 kg (10874 lb)	610 (125)	33 (13)	76.2 (30)	76.2 (30)
		488 (100)	not supported	not supported	not supported
		439 (90)	not supported	not supported	not supported
		342 (70)	not supported	not supported	not supported
Model 986	1099 kg (2421 lb)	610 (125)	13 (5)	76.2 (30)	76.2 (30)
		488 (100)	25 (10)	76.2 (30)	76.2 (30)
		439 (90)	33 (13)	76.2 (30)	76.2 (30)
		342 (70)	61 (24)	76.2 (30)	76.2 (30)
Model 986 and first 86E expansion model	2050.7 kg (4521 lb)	610 (125)	18 (7)	76.2 (30)	76.2 (30)
		488 (100)	43.2 (17)	76.2 (30)	76.2 (30)
		439 (90)	58.5 (23)	76.2 (30)	76.2 (30)
		342 (70)	not supported	not supported	not supported

Table 74. Floor load ratings and required weight-distribution areas (continued)

Configuration _a	Total weight of configuration _b	Floor load rating, kg per m ² (lb per ft ²)	Weight distribution areas _{c, d, e}		
			Sides cm (in.)	Front cm (in.)	Rear cm (in.)
Model 986 with first and second 86E expansion models (with top expansion) ⁶	2985 kg (6581 lb)	610 (125)	21 (8)	76.2 (30)	76.2 (30)
		488 (100)	58.5 (23)	76.2 (30)	76.2 (30)
		439 (90)	not supported	not supported	not supported
		342 (70)	not supported	not supported	not supported
Model 986 with first, second, and third 86E expansion models (with top expansion) ⁶	3919.5 kg (8641 lb)	610 (125)	25 (10)	76.2 (30)	76.2 (30)
		488 (100)	74 (29)	76.2 (30)	76.2 (30)
		439 (90)	not supported	not supported	not supported
		342 (70)	not supported	not supported	not supported
Model 986 with first, second, third, and fourth 86E expansion models (with top expansion) ⁶	4854 kg (10701 lb)	610 (125)	30.5 (12)	76.2 (30)	76.2 (30)
		488 (100)	not supported	not supported	not supported
		439 (90)	not supported	not supported	not supported
		342 (70)	not supported	not supported	not supported
Model 988	1080 kg (2380 lb)	610 (125)	10 (4)	76.2 (30)	76.2 (30)
		488 (100)	21 (8)	76.2 (30)	76.2 (30)
		439 (90)	33 (13)	76.2 (30)	76.2 (30)
		342 (70)	61 (24)	76.2 (30)	76.2 (30)
Model 988 with expansion model 88E	1988 kg (4380 lb)	610 (125)	15.3 (6)	76.2 (30)	76.2 (30)
		488 (100)	38 (15)	76.2 (30)	76.2 (30)
		439 (90)	53.4 (21)	76.2 (30)	76.2 (30)
		342 (70)	not supported	not supported	not supported
Model 988 with first and second expansion model 88E	3050 kg (6722 lb)	610 (125)	23 (9)	76.2 (30)	76.2 (30)
		488 (100)	61 (24)	76.2 (30)	76.2 (30)
		439 (90)	not supported	not supported	not supported
		342 (70)	not supported	not supported	not supported

Table 74. Floor load ratings and required weight-distribution areas (continued)

Configuration _a	Total weight of configuration _b	Floor load rating, kg per m ² (lb per ft ²)	Weight distribution areas _{c, d, e}		
			Sides cm (in.)	Front cm (in.)	Rear cm (in.)
Notes:					
1. A storage system contains a base frame (model 988, 986, 985, 984) and any expansion frames (model 88E, 86E, 85E, 84E) that are associated with it.					
2. The base frame and expansion frame are bolted together. The weights represent a combined configuration, fully populated with all enclosures and adapters.					
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. 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.					
6. The all-flash machine type 533x storage system does not support the configuration noted.					

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 138.
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 dimensions and weights.

Table 75. DS8880 dimensions and weights

Model	Dimensions ¹	Maximum weight
DS8880 Model 984	Height 193 cm (76 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	976 kg (2151 lb)

Table 75. DS8880 dimensions and weights (continued)

Model	Dimensions ¹	Maximum weight
DS8880 Model 985	Height (without top expansion) 193 cm (76 in.) Height (with top expansion) 220 cm (86.5 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	1144 kg (2520 lb)
DS8880 Model 986	Height (without top expansion) 193 cm (76 in.) Height (with top expansion) 220 cm (86.5 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	1099 kg (2421 lb)
DS8880 Model 988	Height (without top expansion) 193 cm (76 in.) Height (with top expansion) 220 cm (86.5 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	1080 kg (2380 lb)
DS8880 Model 84E (first expansion frame) ²	Height 193 cm (76 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	1040 kg (2290 lb)
DS8880 Model 84E (second expansion frame) ²	Height 193 cm (76 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	849 kg (1870 lb)
DS8880 Model 85E (first expansion frame)	Height (without top expansion) 193 cm (76 in.) Height (with top expansion) 220 cm (86.5 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	1116 kg (2460 lb)
DS8880 Model 85E (second and third expansion frames) ²	Height (without top expansion) 193 cm (76 in.) Height (with top expansion) 220 cm (86.5 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	931 kg (2052 lb)

Table 75. DS8880 dimensions and weights (continued)

Model	Dimensions ¹	Maximum weight
DS8880 Model 85E (fourth expansion frame) ²	Height (without top expansion) 193 cm (76 in.) Height (with top expansion) 220 cm (86.5 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	812 kg (1790 lb)
DS8880 Model 86E (first expansion frame)	Height (without top expansion) 193 cm (76 in.) Height (with top expansion) 220 cm (86.5 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	953 kg (2100 lb)
DS8880 Model 86E (second and third expansion frames) ²	Height (without top expansion) 193 cm (76 in.) Height (with top expansion) 220 cm (86.5 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	935 kg (2060 lb)
DS8880 Model 86E (fourth expansion frame) ²	Height (without top expansion) 193 cm (76 in.) Height (with top expansion) 220 cm (86.5 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	935 kg (2060 lb)
DS8880 Model 88E (first expansion frame)	Height (without top expansion) 193 cm (76 in.) Height (with top expansion) 220 cm (86.5 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	908 kg (2000 lb)
DS8880 Model 88E (second expansion frame)	Height (without top expansion) 193 cm (76 in.) Height (with top expansion) 220 cm (86.5 in.) Width 64 cm (25 in.) Depth 144 cm (56.5 in.)	
Notes: 1. These dimensions include casters and covers. The casters are recessed and do not require extra clearance. 2. The all-flash storage system does not support this configuration.		

Service clearance requirements

The service clearance area is the area around the storage system that technical 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, technical service representatives must open the front and rear covers to service the storage system. Figure 18 on page 146 illustrates related service clearance requirements:

- For the rear of the storage system **1**, allow a minimum of 76.2 cm (30.0 in.) for the service clearance.
- The length of the tailgate opening **2** on the storage system.
- For a single storage system installation (from the front of the storage system), allow a minimum of 60.96 cm (24.0 in.) on the right side **3** in the aisle to the front of the machine.
- For the front of the storage system **4**, allow a minimum of 121.9 cm (48.0 in.) for the service clearance.
- The leveling pads **5** and **6** on the storage system.
- The opening of the tailgate **7** on the storage system.

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 151.

Preparing a raised floor for the earthquake resistance kit installation:

You must prepare a raised floor before an earthquake resistance kit can be installed on any frame in your storage system.

About this task

To ensure that you meet all site requirements, obtain the service of a qualified consultant or structural engineer to help you prepare the floor.

Figure 19 and Figure 20 on page 148 illustrate the earthquake resistance kit after it is installed by an IBM service representative on a raised floor.




Figure 19. Turnbuckle assembly




Figure 20. Earthquake resistance kit installed on a raised floor

- 1** Frame
- 2** Support nut
- 3** Leveler feet
- 4** Load plate
- 5** Rubber bushing
- 6** Shaft

Complete the following steps to prepare your raised floor by using Figure 19 on page 147 and Figure 20 as references.

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 frame.

Use Figure 21 as a guide for the location and dimensions of these openings. The pattern repeats for each frame in the configuration. Dimensions are in millimeters (inches).





Figure 21. Locations for the cable cutouts, rubber bushing holes on raised floors, and eyebolt on concrete floors

2. Obtain four fasteners (per frame) that are heavy-duty concrete or slab floor eyebolts. The four fasteners per frame are different than the eight fasteners per frame that are needed for non-raised floor installation. 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 22 on page 150 and Figure 23 on page 151.



f2c02313

- 1** Yoke
- 2** Nut, left-hand thread
- 3** Turnbuckle $\frac{1}{2}$ -13 right hand and left hand
- 4** Nut
- 5** Washer
- 6** Bushing
- 7** Load plates
- 8** Bushing, Plastic
- 9** Rods, $\frac{1}{2}$ -13 right hand, 304 mm, or 533 mm
- 10** Spacer
- 11** Washer flat, thick
- 12** Washer flat, custom

Figure 22. Raised floor tie-down installation




Figure 23. Eyebolt required dimensions. Dimensions are in millimeters (inches).

3. Install the eyebolt fasteners in the concrete floor by using the following guidelines:
 - See Figure 21 on page 149 and Figure 23 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.

About this task

To ensure that you meet all site requirements, obtain the service of a qualified consultant or structural engineer to help you prepare the floor.

Figure 24 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 24. This figure shows two of the most common fasteners that you can use.




Figure 24. 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-pound pull force.
 - The fasteners must have a dimension that fits into the load plate holes, which are each 21 mm (0.826 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. Figure 24 illustrates the 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
2. Work with your consultant or structural engineer and use the following guidelines to install the fasteners in the concrete floor:
- Use Figure 25 on page 154 to determine the placement of the loadplate fasteners **1** (four per loadplate). The pattern repeats for each frame. Dimensions are in millimeters (mm).
 - Ensure that the installed fasteners do not exceed a height of 65 mm (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-pound pull force).
 - If you use a threaded bolt that secures into a threaded insert in the floor and the bolt extends longer than 30 mm (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 25. Location 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 24 on page 152 so that the representative can use these parts secure the load plates onto the floor.

Seismic (earthquake) resistance:

Seismic qualification testing was conducted on the IBM DS8880 Disk Storage Subsystem to meet ICC-ES AC-156 requirements.

Table 76 includes the seismic test parameters and definitions.

Table 76. Seismic test parameters

	S_{DS} (g)	z/h (incl'd roof or lower levels)	HORIZONTAL (g)		VERTICAL (g)		Mounting
			A_{FLX-H}	A_{RIG-H}	A_{FLX-V}	A_{RIG-V}	
ICC-ES AC156, 2012	2.0	1.0	3.20	2.40	1.34	0.54	Raised floor and nonraised floor

Table 76. Seismic test parameters (continued)

	S_{DS} (g)	z/h (incl'd roof or lower levels)	HORIZONTAL (g)		VERTICAL (g)		Mounting
			A_{FLX-H}	A_{RIG-H}	A_{FLX-V}	A_{RIG-V}	
Notes:							
The earthquake resistance kit (feature code 1906) is an optional seismic kit for stabilizing the frame so that the frame complies with IBM® earthquake resistance standards.							
Soil classes A, B, C, D, Seismic Risk Category I, II, III, IV, and Seismic Design Categories A, B, C, D, E, and F are all covered under this certification, limited by the S_{DS} value stated in column 2.							
The following parameters are used in Table 76 on page 154:							
S_{DS}	Design spectral response acceleration at short period						
z/h	Quotient of z (height of the structure with respect to attachment point) and h (average building and structure roof height regarding to the base elevation)						
A_{FLX-H}	Horizontal spectral acceleration calculated for flexible components						
A_{RIG-H}	Horizontal spectral acceleration calculated for rigid components						
A_{FLX-V}	Vertical spectral acceleration calculated for flexible components						
A_{RIG-V}	Vertical spectral acceleration calculated for rigid components						

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 occur, 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 a technical 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 for 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.

DC-UPS circuit breaker switches

Each DC-UPS provides a circuit breaker switch which can be used to remove connection to input line cord and output power of that DC-UPS. If both DC-UPS circuit breakers are switched off all power is removed from the rack immediately.

Attention: These switches are to be used only in specific service actions or in case of emergency. Switching both DC-UPS circuit breakers off in a rack can result in loss of data.

The following power controls can be configured by a technical 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.

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 model.

Important: To eliminate a single point of failure, independent power feeds to each DS8000 power supply are required. At least one of the feeds should have power conditioning to ensure an acceptable level of power quality as specified in standards such as ANSI C84.1 and EN 50160. Further, each power source must have its own wall circuit breaker.

Attention: Ensure that there is a local site disconnect and/or isolation means (such as a service branch circuit breaker, power feed switch, and/or wall socket outlet to an industrial style plug) for each AC main power cord to a system. The internal rack (system side) appliance cord coupler is NOT intended to be connected or disconnected live. Follow local electrical authority site regulations for proper isolation practices such as Lock Out Tag Out (LOTO) and/or approved procedures.

DANGER

Appliance coupler connector is not intended for AC (Alternating Current - Voltage) electrically live plugging nor for interrupting current. Remove and/or isolate input power from the wall end (service branch circuit) to the power cord before attaching or disconnecting the machine end (appliance coupler).



Input voltage requirements

When you plan for the power requirements of the storage system, consider the input voltage requirements.

The following table provides the input voltages and frequencies that the storage system supports.

Table 77. Single-phase input voltages and frequencies

Characteristic	Voltage (single-phase)
Nominal input voltages	200-240 RMS V AC
Minimum tolerated input voltage	180 RMS V AC
Maximum tolerated input voltage	256 RMS V AC
System maximum current rating	DS8886 40 Amps DS8884 24 Amps
Wall breaker rating (1 ph) ¹	DS8886 US, Canada, and regions that require 20% derating factor 50 - 63 Amps DS8886 Geos with no derating requirement 40 - 63 Amps DS8884 All regions 30 - 63 Amps

Table 77. Single-phase input voltages and frequencies (continued)

Characteristic	Voltage (single-phase)
Steady-state input frequencies	50 ± 3 or 60 ± 3.0 Hz
PLD input frequencies (<10 seconds)	50 ± 3 or 60 ± 3.0 Hz
1. Regions such as the United States and Canada require 20% derating for power distribution circuits. For these regions, a minimum 50 A circuit breaker is required for a 40 Amp rated system such as the DS8886. Verification of applicable local standards is required for system installation. Some configuration restrictions to limit power consumption can allow reduced breaker ratings on a regional basis.	

Table 78. Three-phase input voltages and frequencies

Characteristic	Three-phase delta (3ØΔ)	Three-phase wye (3ØY: LL[Line-to-Line])
Nominal input voltages	200, 208, 220, or 240 RMS V AC	380, 400, or 415 RMS V AC
Minimum tolerated input voltage	180 RMS V AC	315 RMS V AC
Maximum tolerated input voltage	256 RMS V AC	456 RMS V AC
System maximum current rating	30 Amps for Japan 24 Amps for other regions	17 Amps
Wall breaker rating (3 ph) ¹	30 - 60 Amps	20 - 32 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
1. Regions such as the United States and Canada require 20% derating for power distribution circuits. For these regions, a minimum 50 A circuit breaker is required for a 40 Amp rated system such as the DS8886. Verification of applicable local standards is required for system installation. Some configuration restrictions to limit power consumption can allow reduced breaker ratings on a regional basis.		

Power connector requirements

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

Table 79 on page 159 shows the power cords and the inline connectors and receptacle 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.

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 DS8880. The DS8880 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.
- 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 79. DS8880 power cords

Power cord feature code	Power cord description ₆	Inline connector	Receptacle	Wire gauge
1062 _{1, 2, 3, 4, 5}	Single-phase power cord, 200V - 240V, 60A, IEC 60309 3-pin customer connector IEC'309 service ready plug for United States, Canada, Latin America, Japan, and Taiwan (EMEA and other regions can optionally order)	HBL360C6W, Pin and Sleeve Connector, IEC 60309, 2P3W	HBL360R6W, ac Receptacle, IEC 60309, 2P3W	10 mm ² (7 AWG)
1063 _{2, 3, 4}	Single-phase power cord, 200V - 240V, 63A, no connector (pre-stripped wires, termination ready for EMEA (except Japan), hard-wired, or uses client provided customer connector/plug as regionally allowed)	Not applicable	Not applicable	10 mm ² (7 AWG)
1086 _{1, 4, 5}	Three-phase WYE (3ØY) voltage (five-wire 3+N+PE), 380-415V~ (nominal line-to-line (LL)), 30 A, IEC 60309 5-pin customer connector IEC'309 service ready plug, for EMEA (other regions under approved circumstances, such as United States, Canada, Latin America, Japan, and Taiwan can optionally order)	HBL530C6V02, Pin and Sleeve Connector, IEC 60309, 4P5W	HBL530R6V02, AC Receptacle, IEC 60309, 4P5W	6 mm ² (10 AWG)
1087 _{1, 4, 5}	Three-phase Delta (3ØΔ) voltage (four-wire 3+PE (Protective Earth Ground)), 200-240V, 30 A, IEC 60309 4-pin customer connector with IEC'309 service ready plug, for United States, Canada, Latin America, Japan, and Taiwan (EMEA and other regions can optionally order)	HBL430C9W, Pin and Sleeve Connector, IEC 309, 3P4W	HBL430C9W, AC Receptacle, IEC 60309, 3P4W	6 mm ² (10 AWG)

Table 79. DS8880 power cords (continued)

Power cord feature code	Power cord description ₆	Inline connector	Receptacle	Wire gauge
1088 ₄	Three-phase WYE (3ØY) voltage (five-wire 3+N+PE), 380-415V~ (nominal line-to-line (LL)), 40 A, no customer connector provided (pre-stripped wires, termination ready for EMEA (except Japan), hard-wired, or uses client provided customer connector/plug as regionally allowed)	Not applicable	Not applicable	6 mm ² (10 AWG)
1089 _{1, 4, 5}	Three-phase Delta (3ØΔ) voltage (four-wire 3+PE (Protective Earth Ground)), 200-240V, 60 A, IEC 60309 4-pin customer connector with IEC 309 service ready plug, for United States, Canada, Latin America, Japan, and Taiwan (EMEA and other regions can optionally order)	HBL460C9W, Pin and Sleeve Connector, IEC 309, 3P4W	HBL460R9W, AC Receptacle, IEC 60309, 3P4W	10 mm ² (8 AWG)

Notes:

1. The customer connector must be IEC 60309-compliant. Power cords that are used with DS8100, DS8300, DS8700, DS8800, and DS8870 cannot be used again with the DS8880.
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.
3. The conductor size for single-phase power cords is 10 mm² (8 awg).
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. The descriptions indicate the line cord rating and are not an indication of system power rating. The cord rating defines the maximum breaker rating for the line cord (60A with connector, 63A without connector). Downstream components are designed for fault currents consistent with facility circuit breaker ratings up to 63A.

The following shows the plug (connector on the line cord).




Figure 26. Single-phase, 200 - 240 V, 60A, IEC60309 cable connector pins




Figure 27. Three-phase low voltage (200 - 240 V) 60A or 30A, IEC60309 cable connector pins




Figure 28. Three-phase high voltage (380 - 415 V) 30A, IEC60309 cable connector pins

The following shows the receptacle or inline connector (the power source, a mirror image of the plug).




Figure 29. Single-phase, 200 - 240 V, 60A, IEC60309 customer outlet pins




Figure 30. Three-phase low voltage (200 - 240 V) phase-to-phase, 60A, IEC60309 customer outlet pins




Figure 31. 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.

Note: The power consumption for the operating disk enclosure is typically 275 watts - 300 watts. The power consumption for the operating High Performance Flash Enclosure Gen2 is 500 watts.

The power consumption and environmental information for the IBM DS8000 is provided in Table 80.

Table 80. Power consumption and environmental information for models 984, 985, 986, 84E, 85E, 86E, 988, 88E

Measurement	Unit of measure	Base frame	Expansion frame
Peak electric power ^{1, 2}	kilovolt amperes (kVA)	Model 984: 4.9; Model 985: 6.7 (single-phase); Model 986: 6.2 (three-phase); Model 988: 8.1 (three-phase)	Model 84E: 4.3; Model 85E: 6.7 (single-phase); Model 86E: 6.4 (three-phase); Model 88E: 4.4 (three-phase)
Thermal load	British thermal units (BTU) per hour	Model 984: 16595; Model 985: 22886 (single-phase); Model 986: 21020 (three-phase); Model 988: 27715 (three-phase);	Model 84E: 14795; Model 85E: 22971 (single-phase); Model 86E: 21743 (three-phase); Model 88E: 15062 (three-phase);
Capacity of exhaust	cubic meters per minute (cubic feet per minute or CFM)	44.2 (1500)	51.8 (1800)

Table 80. Power consumption and environmental information for models 984, 985, 986, 84E, 85E, 86E, 988, 88E (continued)

Measurement	Unit of measure	Base frame	Expansion frame
Ground leakage current	milliamperes (mA)	43	43
Startup current	amperes (A or amp)	≤ 100	≤ 100
Startup current duration	microseconds (μs or μsec)	< 200	< 200
Notes: - The values represent data 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 Flash Enclosures Gen2 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. - All frames and configurations that are used in single-phase mode must not exceed 8 kVA.			

Acoustic declaration

Table 81 describes the acoustic declaration information for the storage system.

Table 81. Acoustic declaration for fully configured version 8.5 storage system

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 984, 985, 986, 988	7.9	7.9	62	61
Model 84E, 85E, 86E, 88E	7.9	7.9	62	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:

- Note where air intake locations are on the models that compose your storage system.
- Verify that you can meet the environmental operating requirements at the air intake locations.
- 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 82 summarizes fan, intake, and exhaust locations.

Table 82. Machine fan location

DS8880 Fan Location	Machine Location	Intake Location	Exhaust Location
Entire machine	Entire	Front covers	Rear covers
Power complex	Bottom	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 83 provides the operating ranges for your storage system when its power is on.

Table 83. 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 84 provides the optimum operating points for your storage system with its power on.

Table 84. Optimum operating points with the power on

Measurement	Value
Temperature	22°C (72°F)
Relative humidity	45%

Table 85 on page 165 provides the operating ranges for a storage system with power on.

Table 85. 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 86 provides the temperatures and humidity requirements for your storage system when the power is off.

Table 86. 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 87 provides the temperatures and humidity requirements for storing your storage system.

Table 87. 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 storage system 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 and shock requirements

The vibration levels that are designed for the storage system comply with class V1 requirements included in the product classes for vibration.

The storage system is designed to operate under the vibration V1 levels that are described in Table 88 on page 166. Additional information includes random vibration PSD profile breakpoints and operational shock levels. See Table 89 on page 166 and Table 90 on page 166.

Table 88. Vibration levels

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 89. Random vibration PSD profile breakpoints

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 90. Operational shock levels

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 Å/month and the rate of weight gain assumes that Cu₂S and Cu₂O grow in equal proportions.
3. The derivation of the equivalence between the rate of silver corrosion product thickness growth in Å/month and the rate of weight gain assumes that Ag₂S 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 technical service representatives. However, the following steps describe the process that is needed to optimize the air circulation and cooling for your storage systems.

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 stand-alone base frames, install two fully perforated tiles in front of each base frame, as shown in Figure 32 on page 168 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 32 on page 168 (**3** 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. An example of a possible minimal arrangement is shown in Figure 32 (4 in the figure).

F = front Base = Base model
R = rear Exp = Expansion model

Single Machine - 2 examples



Hot Aisle / Cold Aisle installation

Figure 32. DS8880 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 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 the storage system can report failures and transmit pertinent debugging data to IBM.

When call home is used, your technical support representative can quickly isolate and resolve potential issues with the storage system. The strategic and preferred remote support connectivity method for call home is internet SSL (Secure Socket Layer) for management-console-to-IBM communication.

For IBM remote access to the management console and the storage system, the preferred remote support connectivity method is Assist On-site (AOS). AOS provides a network-type connection that is secured by SSL and state-of-the-art encryption technology. If your security guidelines allow, AOS can be turned on and configured on the management console, which is the preferred implementation. Alternatively, AOS is installed on a PC that you provide and maintain. Contact your technical support representative for more details.

In addition to AOS, IBM also offers remote support center (RSC). It is a simple SSH proxy based remote service solution that can be used if AOS does not meet your security guidelines. Contact your technical service representative for more information on RSC.

Remote support must be ready at the time of the initial installation of the DS8000. For internet-based remote support, you must open your firewalls.

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.

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 shortwave host adapter Fibre Channel port and the following network components is 380 meters (1,246 ft 9 in.) at 4 Gbps, 150 meters (492 ft 2 in.) at 8 Gbps, and 100 meters (328 ft 1 in.) at 16 Gbps. For a longwave host adapter Fibre Channel port, the maximum distance is 10 km (6.2 miles).
 - 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 technical 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 technical service representatives.

Use the customization worksheets at IBM Techdocs (<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4743>) 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: Technical 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 that technical 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

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, a technical service representative can securely sign on to the management console from a remote location in response to the service call.

The storage system uses secure Internet SSL connectivity for the outbound (call home) remote support connection.

Assist On-Site (AOS) is available 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. AOS can be installed on a customer gateway server in a DMZ or, if your security guidelines allow, it can run directly on the storage system.

In addition to AOS, you can also use 'rsc' (remote support center). It is a simple SSH proxy based remote service solution which can be used if AOS does not meet your security guidelines. Contact your technical service representative for more information on RSC.

Remote code load (RCL) requires outbound (call home) remote support, and either AOS or RSC.

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 customer does not want to send error logs over an encrypted connection.

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 completed any remote service on your storage system, and at what time the work was completed. Contact your technical 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.

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 technical 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

Use the power control 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 technical 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

The following control switch settings are specific to IBM Z.

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 hosts for logging to the Error Recording Data Set (ERDS). SIMs can be selectively reported to the IBM Z 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 LPAR making an I/O request to the logical system or logical volume.

IBM Z high-performance FICON enhanced buffer management

IBM Z 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 storage system vary 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 IBM Z 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 technical support 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 91 compares data migration options and lists advantages and disadvantages of each.

Table 91. 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 91. 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 storage system provides functions to manage data secrecy and networking security, including data encryption, user account management, and functions that enable the storage system to conform with NIST SP 800-131A requirements.

Planning for data encryption

The storage system supports data encryption by using IBM Security Key Lifecycle Manager key servers.

To enable disk encryption, the storage system must be configured to communicate with two or more IBM Security Key Lifecycle Manager key servers. The physical connection between the Hardware Management Console (HMC) and the key server is through an 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

Two encryption-key servers and associated software are required 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 is attached directly 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.

For DS8000 encryption environments a second Hardware Management Console (HMC) should be configured for high availability.

Important: You are responsible for replicating key labels and their associated key material across all key servers that are attached to the encryption-enabled storage system before you configure those key labels on the 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.

For dual-platform key server support, the installation of IBM Security Key Lifecycle Manager interim fix 2 (V1.0.0.2 or later) is recommended to show both key labels in the DS8000 Storage Management GUI. If you intend to replicate keys between separate IBM Z sysplexes by using ICSF with the JCECCARACFKS

keystore in secure-key mode and with the secure-key configuration flag set in IBM Security Key Lifecycle Manager, then IBM Security Key Lifecycle Manager 3 (V1.0.0.3 or later) is required.

To enable encryption on a storage system with version 8.1 (88.10.112.0) or later using TKLM or SKLM, you must upgrade to one of the following versions of TKLM or SKLM that has the Gen2 CA root installed:

- TKLM version 2.0.1 or later on Open Systems
- SKLM (all versions) on Open Systems
- SKLM version 1.1.0.2 or later on z/OS

This SKLM/TKLM upgrade requirement applies to DS8000 shipped with version 8.1 (88.10.112.0) and later.

Planning for key lifecycle managers

DS8000 storage systems support IBM Security Key Lifecycle Manager.

If NIST 800-131A security conformance is required on your storage system, 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 an IBM Z host LPAR with z/OS, use IBM Security Key Lifecycle Manager for z/OS V1.1.0.3 or later.
- If your encryption key server is Gemalto Safenet KeySecure, select version 8.0.0 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 Security Key Lifecycle Manager V2.0.1 or later.
- If your encryption key server runs on an IBM Z host LPAR with z/OS, install IBM Security Key Lifecycle Manager for z/OS v1.0.1 or later.

IBM Storage Appliance 2421 Model AP1 can be ordered either as a single isolated key server (feature code 1761) or as two isolated key servers (feature codes 1761 and 1762, ordered together). This order must include an indicator for IBM Security Key Lifecycle Manager (feature code 0204), which indicates that a DVD with IBM Security Key Lifecycle Manager software is provided with Storage Appliance AP1. For more information, search for "IBM Storage Appliance 2421 Model AP1" at the IBM Publications Center website (www.ibm.com/shop/publications/order).

If you want to acquire a different isolated key server, refer to the *IBM Security Key Lifecycle Manager Installation and Configuration Guide* (SC27-5335) or IBM Security Key Lifecycle Manager online product documentation (www.ibm.com/support/knowledgecenter/SSWPVP/) for hardware and operating system requirements.

Note: You must acquire an IBM Security Key Lifecycle Manager license for use of the IBM Security Key Lifecycle Manager software that is ordered separately from the stand-alone server hardware. The IBM Security Key Lifecycle Manager license

includes both an installation license for the IBM Security Key Lifecycle Manager management software and a license for encrypting drives.

IBM Security Key Lifecycle Manager for z/OS generates encryption keys and manages their transfer to and from devices in an IBM Z environment.

Planning for full-disk encryption activation

Full-disk-encryption drives are standard on the storage system. 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 us**.

You are responsible for downloading or obtaining from IBM, and installing designated machine code (such as microcode, basic input/output system code [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.
5. For applications that require network access to the storage system, assign a unique user ID (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, storage systems allow 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 DS8880 storage system requires 112-bit security strength on all SSL/TLS connections, other than remote support 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 storage systems, you must ensure that all external system networks connected to the DS8880 storage systems 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 (http://www.ibm.com/support/knowledgecenter/ST5GLJ_8.1.0/com.ibm.storage.ssic.help.doc/f2c_securitybp.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.

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 technical 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 the DSFA website, 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 page 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 the storage system.

The Base Function license is based on the entire capacity of the storage system.

The z-synergy Services, Copy Services, and Copy Services Manager on the Hardware Management Console (HMC) licenses are priced based on the use of license capacity that you purchase.

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 Services 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 functions 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 function. You are notified by IBM and required to extend enablement of your licensed function and install a new licensed feature key.

Activation of licensed functions

After the technical 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 feature activation codes to your storage system. You can apply the feature 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: Removal of a DS8000 licensed function to deactivate the function is non-disruptive, but takes effect at the next IML.

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:

- A removable media for downloading your feature 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 feature activation 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 a technical 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 function, such as Copy Services.

About this task

For example, assume that you initially purchased 250 TB of Copy Services capacity. After several months, you need an extra 100 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. For example, you order four of the Copy Services feature code 8253 (25 TB each for a total of 100 TB) against the serial number of the 533x or 904x machine type license currently on your storage system. This additional license capacity increases your Copy Services 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 the feature activation code for the licensed feature. This new feature activation code represents the total capacity that is now licensed (or 350 TB). It licenses the original 250 TB plus the additional 100 TB that you just ordered.
4. Obtain the feature activation code for the licensed feature from your sales representative.
5. After you obtain the feature activation code for the licensed feature, enter it into the DS8000 Storage Management GUI. You replace the existing feature activation code with the new feature activation code.
6. After the feature activation code is installed successfully, you now have 350 TB of Copy Services 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 function to make that licensed function inactive.

About this task

If you have an active optional function and you want to replace it with an inactive function, you must repurchase the function 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 function, up to the previously purchased level, without having to repurchase the function.

Appendix A. Accessibility features

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 information on machine type models 283x or 533x.

Appendix C. IBM equipment and documents

Use the documents provided with your storage system 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 your technical 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
- Expansion frames

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 primary management console is included with each base frame that you order. The management console is physically located (installed) inside the base frame. If needed, you can order a secondary management console for the base frame.

Power cords

Your shipment includes the power cords that you ordered.

Various media

IBM ships the following media, 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, 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 283x or 533x 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 technical 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 technical 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 technical service representatives can use them when they service your storage system.

Service media

Your delivery includes the following media for technical service representatives to use:

- Operating system media
- Management console media:
 - Management console critical backup SDHC memory card
 - Dump, trace, statesave SDHC memory card, which technical support 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. DS8870 to DS8880 model conversion

You can convert an existing DS8870 (models 961 and 96E) to a high-performance DS8886 (machine type 283x models 985 and 85E or 986 and 86E).

DS8870 to DS8880 model conversion summary

The conversion is a non-concurrent operation requiring your planning and several days of onsite IBM support with your full participation.

Model conversion summary:

- Unlike previous model conversions that only changed the model numbers, this conversion also changes the machine type to 2831. This can require alterations to your machine asset depreciation scheme.
- The only machine type for the model conversion is the 2831 with a one year service warranty. The existing DS8870 warranty period is not applied.
- This model conversion changes the models from 961 and 96E to machine type 283x model 985 and 85E or 986 and 86E.
- Each DS8870 that is converted to DS8880 retains the existing (original DS8870) frame serial number and worldwide node name (WWNN).
- Some configurations require more DS8880 frames to contain the same amount of data in a DS8870 configuration. Each additional frame has a unique new serial number. To determine the number of DS8880 frames required for your configuration, contact IBM.
- IBM provides new frames, including the management console(s), processor enclosures, power supplies, and I/O enclosures. The first expansion frame includes power supplies and I/O enclosures. The second, third, and fourth expansion frames include power supplies.

Note: The High Performance Flash Enclosures are not transferred to the new 985, 986, 85E, and 86E models.

- The device adapters, host adapters, storage enclosure assemblies, and drives from the existing DS8870 frames are moved into the new DS8880 frames.
- Each DS8870 frame (without the adapters, storage enclosures, and disk drives) is returned to IBM.
- Apply the DS8880 licensed machine-code keys after conversion and before logical configuration is created. The DS8880 has a unique license feature. You cannot apply the licensed machine-code keys from the DS8870.
- Any additional upgrades that are planned for DS8880 must be done after the conversion is completed.

You must complete all of the prerequisites in the following table before the IBM service representative starts the model conversion:

Table 92. Model conversion prerequisites

Prerequisites	Description
Data migration, logical configuration removal, encryption group removal	You need to migrate all data, remove all logical configuration, and remove the encryption group (if encryption is activated), from the DS8870.

Table 92. Model conversion prerequisites (continued)

Prerequisites	Description
Secure Data Overwrite	Secure Data Overwrite (SDO) is optional. DS8870 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 perform an SDO. This is an extra service. The SDO 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 data drives. This process can take significant time to complete. You need to plan accordingly.
Data encryption and Security Key Lifecycle Manager servers	If the DS8870 uses data encryption and Security Key Lifecycle Manager servers and you intend to activate encryption on the DS8880, ensure that the infrastructure is in place for the DS8880 system. Encryption is optional for conversion.
Ethernet cable and Fibre Channel host cables	Order the following cables for the new DS8880 only if your existing DS8870 infrastructure cannot be used: • Fibre Channel host cables from the host to DS8880 host adapters. • An Ethernet LAN signal cable. This cable is required to connect the management console to your LAN. If the DS8870 configuration has an external secondary management console, then the new DS8880 will have a secondary management console installed inside the base frame. You might need to plan to move the support infrastructure. Note: If you want the IBM service representative to route the cables, that activity must be coded or billed separately from the DS8880 installation work.
Power cables	DS8880 requires a new set of power cords and new set of perforated floor tiles. Depending on your configuration, it might also require a different wall breaker. For more information, see “Planning for power requirements” on page 155.
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 DS8880 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 DS8870 cannot be reused.

Table 92. Model conversion prerequisites (continued)

Prerequisites	Description
Physical installation requirements	Determine whether the new DS8880 can occupy the same floor space as the existing DS8870. • If you are using the same location, you need to provide a location for both the existing DS8870 and converted DS8880 frames so that the hardware can be staged between them. • To determine the number of DS8880 frames required for your configuration, contact IBM. If your configuration requires additional frames, ensure that you have sufficient physical space. • Whether installed in an existing or new location, the location must conform to the requirements for DS8880.
Management console settings	Determine how the IBM service representative receives the settings for the new management consoles (primary and secondary). If the DS8870 configuration has an external secondary management console, then the new DS8880 will have a secondary management console installed inside the base frame. You might need to plan to move the support infrastructure. • 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	DS8880 does not support a 2x2 configuration. The DS8870 must be in a 1x1 configuration prior to conversion, and must not be connected to another storage system by the black and gray private network cables. If the DS8870 is not already in a 1x1 configuration, the IBM service representative must use discontinue procedures to separate it before you begin the conversion. If the existing DS8870 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.
Serviceable events	Before you begin the conversion, the IBM service representative must repair any open serviceable events. Hardware repairs add service time.

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 can differ from the previous model. For more information, see “Feature codes for power cords” on page 110.
3. Verify that the circuit breaker rating meets the storage system requirements. DS8880 power requirements are completely different from the DS8870 requirements. For more information, see “Input voltage requirements” on page 157
4. Verify that the physical requirements necessary for 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 technical 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 technical 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, Security Key Lifecycle Manager, 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 technical service representative to perform a Secure Data Overwrite (SDO). This is an extra service. 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 the DS8880:

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 obtained from the IBM DSFA website.
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. New features for models 980, 981, 982, 98B, 98E, and 98F

For models 980, 981, 982, 98B, 98E, and 98F, support is provided for High Performance Flash Enclosures Gen2 (including the 3.8 TB and 7.6 TB Flash Tier 1 drives and Flash Tier 2 drives), zHyperLink, and Transparent cloud tiering.

Note: Models 980, 981, 982, 98B, 98E, and 98F are supported, but not covered in this version documentation. For more information on these models, refer to previous version documentation (GC27-8525-06).

Table 93. License Machine Code feature codes to update models 980, 981, 982, 98B, 98E, and 98F

Feature code	Feature	Description
1883	DS8000 Licensed Machine Code R8.3	Microcode bundle 88.x.xx.x for base models 980, 981, and 982
1884	DS8000 Licensed Machine Code R8.4	Microcode bundle 88.x.xx.x for base models 980, 981, and 982 Required for 7.6 TB Flash Tier 2 drive support
1983	DS8000 Licensed Machine Code R8.3	Microcode bundle 88.x.xx.x for expansion models 98B, 98E, and 98F
1984	DS8000 Licensed Machine Code R8.4	Microcode bundle 88.x.xx.x for expansion models 98B, 98E, and 98F Required for 7.6 TB Flash Tier 2 drive support

Support for High Performance Flash Enclosures Gen2

Placement of a High Performance Flash Enclosure Gen2 pair at the first available (starting from the bottom) standard enclosure pair location within the system is supported for models 980, 981, 982, 98B, 98E, and 98F.

The High Performance Flash Enclosure Gen2 pair provides two 2U storage enclosures with associated flash RAID controllers and cabling. This combination of components forms a high-performance, fully-redundant flash storage array. The array components can be exchanged without interruption of service for concurrent maintenance and firmware updates.

- Models 980 and 98B each support up to one High Performance Flash Enclosure Gen2 pair
- Models 981 and 98E each support up to two High Performance Flash Enclosure Gen2 pair
- Model 982 supports up to two High Performance Flash Enclosure Gen2 pair, and Model 98F supports up to four High Performance Flash Enclosure Gen2 pair

Table 94. High Performance Flash Enclosure Gen2 feature codes for models 980, 981, 982, 98B, 98E, and 98F

Feature code	Feature	Description
1303	Gen2 I/O enclosure pair	High Performance Flash Enclosure Gen2 support for models 980, 981, 982, 98B, 98E, and 98F Required to support feature codes 1600, 1601, and 1604
1601	Flash RAID controller pair	High Performance Flash Enclosure Gen2 support for models 980, 98B, 981, 98E Requires feature code 1602
1602	High Performance Flash Enclosure Gen2 pair	Requires feature code 1601 or 1604
1604	Flash RAID adapter pair	High Performance Flash Enclosure Gen2 support for models 982 and 98F Requires feature code 1602
1610	400 GB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives)
1611	800 GB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives)
1612	1.6 TB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives)
1613	3.2 TB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives)
1623	3.8 TB 2.5-inch Flash Tier 1 drive set	Flash drive set (16 drives) No intermix with Flash Tier 0 or Flash Tier 2 drive sets
1624	7.6 TB 2.5-inch Flash Tier 2 drive set	Flash drive set (16 drives) No intermix with Flash Tier 0 or Flash Tier 1 drive sets
1699	High Performance Flash Enclosure Gen2 filler set	Includes 16 fillers

zHyperLink

For storage systems with 12-core and higher processors, zHyperLink connections with z Systems hosts provide low latency for random reads.

The z-synergy Services license is required for zHyperLink. I/O Priority Manager must be disabled.

The following table lists feature codes for zHyperLink.

Table 95. zHyperLink feature codes for storage systems with 12-core and higher processors

Feature code	Feature	Description
1303	Gen2 I/O enclosure pair	High Performance Flash Enclosure Gen2 support for models 980, 981, 982, 98B, 98E, and 98F Required to support feature code 3500
1450	zHyperLink cable	40 m (131 ft), OM4 50/125 micrometer, multi mode, MTP connectors Requires feature code 3500 For storage systems with 12-core and higher processors
1451	zHyperLink cable	150 m (492 ft), OM4 50/125 micrometer, multi mode, MTP connectors Requires feature code 3500 For storage systems with 12-core and higher processors
3500	zHyperLink I/O-adapter	For storage systems with 12-core and higher processors Required for feature code 1450 and 1451
4235	256 GB system memory	12-core For model 980
4425	12-core POWER8 processors	Requires feature code 4235 For model 980

Transparent cloud tiering adapters

Use these feature codes to order adapter pairs to enhance Transparent cloud tiering connectivity for your storage system.

Transparent cloud tiering connectivity can be enhanced with 10 Gbps adapter pairs to improve bandwidth for a native cloud storage tier in IBM Z environments.

Table 96. Feature codes for Transparent cloud tiering adapter pairs

Feature code	Description	Models
3600	2-port 10 Gbps SFP+ optical/2-port 1 Gbps RJ-45 copper longwave adapter pair for 2U processor complex	model 980
3601	2-port 10 Gbps SFP+ optical/2-port 1 Gbps RJ-45 copper longwave adapter pair for 4U processor complex	models 981 and 982

Appendix F. Customization worksheets

You must complete the customization worksheets before any installation of a new storage system or management console.

The customization worksheets are frequently updated. To ensure that you use the latest version, download a spreadsheet file that contains all of the worksheets from IBM Techdocs (<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4743>).

Use the 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, power control, and control switch settings. For more details, refer to Chapter 7, “Planning your storage complex setup,” on page 173.

After you fill out the worksheets, give them to the technical service representatives who complete the installation.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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.

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 actual people or business enterprises 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.

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” on page 219 and the “Caution notices” 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

Ensure that you understand the caution notices.

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: This part or unit is heavy but has a weight smaller than 18 kg (39.7 lb). Use care when lifting, removing, or installing this part or unit. (C008)

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 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:

Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. Although shining light into one end and looking into the other end of a disconnected optical fiber to verify the continuity of optic fibers may not injure the eye, this procedure is potentially dangerous. Therefore, verifying the continuity of optical fibers by shining light into one end and looking at the other end is not recommended. To verify continuity of a fiber optic cable, use an optical light source and power meter. (C027)

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


CAUTION:



or



or



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: Activate locks or brakes, or apply chocks as directed. Parts may shift or fall and cause personal injury or mechanical damage if these safeguards are not used. (C042)

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)

CAUTION: The rack cabinet is supplied with native built-in extendable outriggers with small floating supplemental castors as motion anti-tip features. They must all be extended into a latched position before and during cabinet movement or relocation. These native built-in outriggers must not be removed completely, but rather recessed in when finished to ensure they are readily available for future use. (C050)

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

Ensure that you understand the danger notices.

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:

- If IBM supplied a power cord(s), 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.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements.
- Do not continue with the inspection if any unsafe conditions are present.
- 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)

Labels
Labels

Use the reference numbers in parentheses at the end of each notice, such as (L001), to find the matching translated notice in *IBM Systems Safety Notices*.

DANGER: Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)

DANGER: Multiple power cords. The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. (L003)

CAUTION: System or part is heavy. The label is accompanied by a specific weight range. (L009)

CAUTION: Pinch hazard. (L012)

ATTENTION: Refer to the Installation Planning Guide for additional explanation on configurations supported for given voltage groups and Power Cord information. (L025)

ATTENTION: For use at altitude 2000 m or lower (L026)

Text based labels
Text Based Labels

 DANGER HIGH LEAKAGE CURRENT Earth connection essential before connecting supply FARE OVERLEDNING Alt udsat mid kobles til jordede stikkontakter PERICO ALTA CORRENTE DE DISPENSÃO É essencial fazer ligação terra antes de fazerem a ligação FARA LÄCKSTROM Enheten måste vara jordad innan den ansluts NEBEZPEČENOSTVO MOŽNOST PRENIKNI NAPÁVIA Uzemní pout před propojení zápisu	OPASNOST JAKO PROBLÉME STRUJE Uzemljenje je nužno prije spajanja dovode napajanja NEBEZPEČI VYSOKÝ ÚNIKÁČÍ PRŮUD Před propojením zařízení je nutnost připojit zemní vodič Fare Hög lekström Tillslutning till jord är nödvändig, för att sluten tillslutas GEVAAR HOOG LEKSTROOM Eerst aardeverbinding maken, daarna voeding aansluiten VORSICHT HOHER ABLEITSTROM Vor Anschließen der Stromversorgung unbedingt erden	PERICOLO CORRENTE DI DISPERSIONE ELEVATA È essenziale effettuare il collegamento di terra e proteggere prima di collegare l'alimentazione PERIGRO CORRIENTE DE ALTA DISPERSION Es imprescindible la conexión a tierra antes de conectar el suministro VARO KORKEA VUOTOVIRTA Suojamaailmalla on varmistettava, ennen kuin laita liitetään sähkökaapeliin خطر فارسه الأرضية ضرورية قبل توصيل التيار KLINGENS YUNDO PEYMA (BAPKHE) Anatoma, yelutun qur kullu kullu	DANGER COURANT DE FUITE IMPORTANT Mettre à la terre obligatoirement avant la reconnexion au réseau d'alimentation principale Figyelem! NAGY SZIVÁRGÓ ÁRAM Földelés feltétlenül szükséges az áramellátás előtt PERICOLO ELEVADA CORRENTE DE FUGA Imprescindible a ligadura a terra antes de conectar do alimentador OTACHO BACOKKIR TOK YTEKIR Před propojením zařízení nutno připojit podzemní NEBEZPEČENOSTVO ZWISZCZ PRĄD ULEWNY Przed podłączeniem urządzenia należy skonektować podziemną	危険 高漏れ電流 接地を済ませ、必ず接続。 危險 高漏電流 連接前，務必先完成接地。 危險 高漏電流 電源を接続する前に必ず接地してください。 注意 누전 위험 전원 공급전해 반드시 먼저 연결함. סכנה זרם גרירה גבוה חיבור למארקה חובה לפני חיבור למערכת הזרם
--	--	---	---	---

and / or



Environmental notices

This information contains all the required environmental notices for IBM Systems products in English and other languages.

The *IBM Systems Environmental Notices* (ftp://public.dhe.ibm.com/systems/support/warranty/envnotices/environmental_notices_and_user_guide.pdf) information includes statements on limitations, product information, product recycling and disposal, battery information, flat panel display, refrigeration and water-cooling systems, external power supplies, and safety data sheets.

Electromagnetic compatibility notices

The following Class A statements apply to IBM products and their features unless designated as electromagnetic compatibility (EMC) Class B in the feature information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices that are supplied with the monitor.

Canada Notice

CAN ICES-3 (A)/NMB-3(A)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Warning: This equipment is compliant with Class A of CISPR 32. In a residential environment this equipment may cause radio interference.

Germany Notice

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

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55032 Klasse A ein.

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

von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden: "Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)." Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse A

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

Verantwortlich für die Einhaltung der EMV-Vorschriften ist der Hersteller:

International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:

IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 800 225 5426
e-mail: Halloibm@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55032 Klasse A.

Japan Electronics and Information Technology Industries Association (JEITA) Notice

(一社) 電子情報技術産業協会 高調波電流抑制対策実施 要領に基づく定格入力電力値: Knowledge Centerの各製品の 仕様ページ参照
--

This statement applies to products less than or equal to 20 A per phase.

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類：6（単相、P F C回路付）
- 換算係数：0

This statement applies to products greater than 20 A per phase, three-phase.

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類：5（3相、P F C回路付）
- 換算係数：0

Japan Voluntary Control Council for Interference (VCCI) Notice

この装置は、クラスA情報技術装置です。この装置を家庭環境で使用すると電磁妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

VCCI-A

Korea Notice

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서 가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

People's Republic of China Notice

声 明

此为 A 级产品,在生活环境中,
该产品可能会造成无线电干扰。
在这种情况下,可能需要用户对其
干扰采取切实可行的措施。

Russia Notice

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

nusemi

Taiwan Notice

警告使用者：

這是甲類的資訊產品，在
居住的環境中使用時，可
能會造成射頻干擾，在這
種情況下，使用者會被要
求採取某些適當的對策。

taiemi

IBM Taiwan Contact Information:

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

f2c00790

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

(1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Index

Numerics

- 2.5-inch 47
- 3.5-inch 47
- 4-port HA 48
- 8-port HA 48

A

- accessibility
 - features 193
- acclimation 127
- activating licenses 190
- adapters 170
 - configuration rules 106
 - Fibre Channel host 99, 101
 - zHyperLink 101
- air circulation 163, 167
 - intake and exhaust 164
- air quality 165
- algorithms 53
- allocation methods 34
- AS/400 LUN
 - control switch settings 175
- Attachment to IBM Z 175
- auto-rebalance 61
- auxiliary volumes 61
- availability features 46

B

- balancing the I/O load 50
- Base function license
 - Database Protection 119
 - Easy Tier 120
 - Encryption Authorization 120
 - I/O Priority Manager 120
 - Operating environment license 120
- Base Function license 118
 - thin provisioning 120
- battery assemblies 112
- battery service modules
 - feature codes 112
- BSMI certificate 113, 114
- BTU 162

C

- cable
 - configuration rules 106
 - cutout locations 135
 - disk drive 89
 - drive cables, feature codes 90
 - feature codes, Fibre Channel
 - cable 101
 - feature codes, zHyperLink cable 102
 - Fibre Channel host adapter 99
 - I/O adapters 99
 - I/O cables 99
 - installation 135

- cable (*continued*)
 - overhead cable 135
 - RIO-G 99
 - top exit bracket 135
 - zHyperLink adapter 101
- cables
 - drive 91
- cache 109
- capacity
 - calculating physical and effective 93
 - exhaust 162
 - floor load rating 138
- caution notices 213, 214
- CCW, channel command words 51
- certificate, BSMI 113, 114
- circuit breakers
 - high-voltage 158
 - low-voltage 158
- CKD, count key data storage 51
- clearances required for service 145
- CLI, command-line interface 35
- cloud
 - transparent cloud tiering 57
- clusters, RAID disk groups 50
- CoD 47
- cold weather procedures 127
- comments, sending x
- communication requirements, host
 - attachment 170
- company information 173
- configuration
 - battery service modules 112
 - disk drive cables 89
 - I/O (RIO-G) cables 99
 - processor memory 109
- configuration control indicators 83
- configuration overview 27
- Configuration overview 29
- configuration rules
 - device adapters 103
 - flash interface cards 103
 - host adapters and cables 106
 - I/O adapter 103
 - I/O enclosures 103
 - management consoles 85
 - Standby CoD disk drives 91
 - storage devices 91
 - storage enclosures 91
- connectivity
 - I/O enclosures 99
- consolidating storage 50
- containers, shipping 128
- contamination information 166
- control switch settings 175
- control unit threshold 175
- conventions
 - terminology v
 - typefaces v
- cooling 163, 167
- Copy Services 37
 - considerations 67

- Copy Services (*continued*)
 - disaster recovery 76
 - licensed functions 80
 - overview 67
 - Safeguarded copy function 124
- Copy Services license 123
 - FlashCopy function 124
 - Remote mirror and copy 124
 - z/OS Global Mirror 125
 - z/OS Metro/Global Mirror
 - Incremental Resync 125
- Copy Services Manager 37
- Copy Services Manager on hardware
 - management console license 125
- Copy Services software package
 - Remote mirror and copy 124
- corrosive gasses and particulates 165
- count key data storage 51
- CUIR, control-unit initiated
 - reconfiguration 175
- customization worksheets 209

D

- danger notices 213, 219
- data
 - securing 81
- data migration
 - selecting method 180
- data placement 50
- Database Protection 119
- DC-UPS 112
- description
 - VMware 64
- description of Easy Tier 61
- description of EAV 54
- device adapters 89
 - configuration rules 103
 - feature codes 89
- Device threshold 175
- dimensions
 - storage system, installed 142
- disaster recovery
 - Copy Services 76
- disk drive
 - cable 85
 - cables 89
- disk drive module
 - maintenance policy 47
- disk drive sets 85
- disk drives 47
- disk enclosures 85
 - fillers 85
- disk intermix
 - configuration indicators 83
- Disk Manager
 - monitoring performance 37
- Disk Storage Feature Activation (DSFA) 190
- drive enclosures 88

- drive set
 - capacity 93
- drive sets 86
- drives
 - cable 91
 - capacity calculation 93
- DS command-line interface 35
- DS8000 Storage Management GUI 35
- DS8870 199
- DS8870 to DS8880
 - model conversion 199
- DS8880 199
- DSFA, Disk Storage Feature
 - Activation 190
- dynamic expansion
 - volume 59
- dynamic volume expansion 59

E

- EAM 52
- earthquake preparedness 169
- earthquake resistance kit
 - required preparation for 146, 152
- Easy Tier 34, 120
 - overview 61
- EAV CKD
 - 1 TB IBM Z CKD 54
 - 3390 Model A 54
 - cylinder size 54
- enclosure fillers
 - feature codes 89
- encryption
 - overview 81
 - planning for 183, 184
- Encryption Authorization 120
- environment 162
 - air circulation 164
 - operating requirements 163, 164
- environmental notices 222
- ePLD
 - See* extended power-line disturbance
- EPLD
 - See* extended power line disturbance
- ESE capacity controls 60
- ESE volumes 61
- exhaust 162
- expansion model position
 - configuration indicators 83
- extended address volumes
 - overview 54
- extended power line disturbance 113
- extended power-line disturbance 113
- Extended Remote Copy (XRC) (see z/OS Global Mirror) 125

F

- failover and fallback 76
- FB, fixed block 51
- FDE 47
- Feature Activation (DSFA), IBM Disk
 - Storage 190
- feature codes
 - See also* features
 - additional setup options 113

- feature codes (*continued*)
 - battery service modules 112
 - device adapters 89
 - drive cables 90
 - drive enclosures 88
 - drive sets 86
 - enclosure fillers 89
 - extended power line disturbance 113
 - extended power-line disturbance 113
 - features, other configuration 113
 - Fibre Channel cable 101
 - Fibre Channel host adapters 100
 - flash drives 86
 - flash RAID adapters 102
 - I/O (RIO-G) cables 99
 - I/O adapter 98
 - I/O enclosures 99
 - management console 85
 - memory 109
 - optional 113
 - ordering optional features 83
 - overhead cable management 134
 - physical configuration 83
 - power cords 110
 - power features 113
 - processors 109
 - setup 113
 - shipping 113
 - shipping and setup 113
 - shipping weight reduction 114
 - Transparent cloud tiering
 - adapters 102, 207
 - zHyperLink cable 102
 - zHyperLink I/O adapters 101
- features
 - input voltage
 - about 112
- feedback, sending x
- Fibre Channel
 - host adapters 99
 - host adapters feature codes 100
 - host attachment 170
 - open-systems hosts 49
 - SAN connections 49
- Fibre Channel cable
 - feature codes 101
- fire suppression 169
- fixed block storage 51
- flash copy 78
- flash drives 47, 86
- flash drives, feature codes 86
- flash interface cards
 - configuration rules 103
- flash RAID
 - adapters feature codes 102
- flash RAID adapters 89
- FlashCopy
 - Multiple incremental FlashCopy 70
- floor and space requirements 131
- floor load 138
- force option 52, 59

G

- Global mirror 124
- Global Mirror 76, 78

H

- HA intermix 48
- hardware features 41
- hardware planning 41
- High Performance FICON 122
- High Performance Flash Enclosure Gen2
 - for previous models 205
- high-voltage
 - installations 158
- homogeneous pools 61
- homologation 213
- host adapters
 - Fibre Channel 99
- host attachment
 - overview 48
- host systems
 - communication requirements 170
- hot spot management 61
- how to order using feature codes 83
- HyperPAV 122

I

- I/O adapter
 - configuration rules 103
 - features 98
- I/O cable
 - configuration 99
- I/O enclosure 48
- I/O enclosures 99
 - configuration rules 103
 - feature codes 99
- I/O load, balancing 50
- I/O plug order 48
- I/O Priority Manager 120
- IBM Copy Services Manager
 - copy services 37
 - Replication 37
- IBM Disk Storage Feature Activation (DSFA) 190
- IBM Spectrum Control 37
- IBM Z
 - HyperPAV 122
 - parallel access volume 122
 - Parallel Access Volumes 122
 - power control settings 175
- IBM Z hosts
 - FICON attachment overview 49
- IEC 60950-1 213
- implementation, RAID 31
- initialization 55
- input voltage 112
 - configuration rules 113
- input voltage requirements 157
- installation
 - air circulation 167
 - components with shipment 197
 - nonraised floor 131
 - raised floor 131
- installation planning 127
- IOPS 61

L

- labels 221
- labels, safety information 213

- LFF 47
- licensed functions 190
 - Copy Services 80
- licenses
 - Disk Storage Feature Activation (DSFA) 190
 - function authorization documents 190
- logical subsystems
 - overview 33
- logical volumes 52
- low-voltage
 - installations 158
- LSS, logical subsystem 33
- LUN
 - calculation 53
 - control switch settings 175

M

- machine types 4
- maintenance policy
 - disk drive module 47
- managed allocation 34
- management console
 - configuration rules 85
 - ESSNI 84
 - feature codes 85
 - HMC 84
 - HMC remote access 84
 - multiple storage systems 46
 - network settings 173
 - overview 46
 - TCP/IP ports 170
- management consoles
 - feature codes 84
- management interfaces 35
- Media threshold 175
- memory
 - feature codes 109
 - system memory 109
- Metro mirror 124
- Metro/Global mirror 124
- migrating data
 - selecting method 180
- model conversion 199
 - DS8870 to DS8880 199
- Multiple incremental FlashCopy 70

N

- network settings 173
- new features xi
- nodes 170
- noise level 162
- notices
 - caution 214
 - danger 219
 - environmental 222
 - labels 221
 - safety 213
- notification settings
 - methods 175

O

- obtaining activation codes 190
- operating environment
 - power on 164
 - while in storage 164
 - with power on or off 164
- Operating environment license 120
- overview
 - host attachment 48

P

- Parallel Access Volumes 122
 - understanding static and dynamic 66
- pass-through 78
- PAV (Parallel Access Volumes) 66
- performance gathering
 - Disk Manager 37
- physical configuration
 - drive capacity 93
 - drive enclosures 86, 88
 - drives 86
 - extended power line disturbance 113
 - I/O adapter features 98
 - I/O cable 99
 - I/O enclosures 99
 - input voltage of power supply 112
 - keyboard 85
 - management console 85
 - management consoles 84
 - power cords 110
 - power features 113
 - processors 109
- planning
 - activating
 - full-disk encryption 185
 - disk encryption
 - activating 185
 - planning 185
 - earthquake resistance kit site preparation 146, 147, 152
 - encryption 183, 184
 - environment requirements 163
 - floor load 138
 - full-disk encryption activation 185
 - IBM Security Key Lifecycle Manager 183, 184
 - IBM Storage Appliance 2421 Model AP1 184
 - isolated key servers 184
 - model conversion 199
 - network and communications 169
 - operating environment, power 165
 - power connector 158
 - safety 169
 - storage complex setup 173
 - weight 138
- point-in-time copy 124
- pool rebalancing 61
- power
 - consumption 162
 - extended power line disturbance feature 113
 - operating environment, off 165
 - operating environment, on 164
 - outlet requirements 156

- power connector
 - requirements 158
 - specifications 158
- power cords 110
 - feature codes 110
- power connector 158
- power features
 - configuration rules 113
- power frequencies 157
- power supply
 - input voltage of 112
- power control settings 175
- Preparing a raised floor for the earthquake resistance kit installation
 - earthquake resistance kit required preparation for 147
- Present SIM data to all hosts 175
- processor
 - feature codes 109
 - memory (cache) 109
- processors
 - feature codes 109
- publications
 - ordering x
 - product v
 - related v

Q

- quick initialization 55

R

- RAID
 - disk groups 50
 - implementation 31
 - RAID 10 overview 32
 - RAID 5 overview 32
 - RAID 6 overview 32
 - RAID overview 31
- raised floors
 - cutting tiles for cables 135
- rank depopulation 61
- redundancy
 - management consoles 84
- Remote mirror and copy 124
- remote mirror for z Systems (see z/OS Global Mirror) 125
- remote power control 170
- remote support
 - connections 170
 - settings 174
- replication
 - copy services functions 37
- requirements
 - floor and space 131
 - floor load 138, 142
 - host attachment communication 170
 - input voltage 157
 - loading dock 129
 - planning network and communications 169
 - power connectors 158
 - power outlets 156
 - receiving area 129
 - service clearance 145

- requirements (*continued*)
 - space 142
- resource groups
 - copy services 78
- RGL 78
- RIO-G cable 99
- rotate capacity 34
- rotate volumes 34

S

- safety 169
 - earthquake preparedness 169
 - earthquake resistance kit 146, 147, 152
 - fire suppression 169
 - information labels 213
 - notices 213
 - operating environment 169
 - power outlets 156
 - temperature and cooling 169
- SAN
 - connections with Fibre Channel adapters 49
- SAS 47
- SAS enterprise and NL SAS 61
- SATA 47
- scenarios
 - adding storage 191
- scope limiting
 - disaster recovery 78
- security
 - best practices
 - service accounts 186
 - user accounts 185
- serial number setting 175
- service clearance requirements 145
- Service Information Messages 175
- SFF 47
- shipments
 - authorized service components 198
 - container weight, dimensions 128
 - hardware, software 197
 - loading ramp 129
 - media 198
 - planning for receipt 129
 - receiving area 129
 - reducing weight of 114
 - requirements
 - loading ramp 129
 - weight reduction feature code 114
- shipping containers 128
- shipping weight reduction 114
 - feature code 114
- SIM 175
- slot plug order 48
- space requirements 142
- specifications
 - power connectors 158
- standard drive enclosures 86
- standards
 - air quality 165
- Standby CoD disk drives 91
- statement of limited warranty 195
- storage area network
 - connections with Fibre Channel adapters 49

- storage features
 - configuration rules 91
 - drives and enclosures 86
- storage image
 - cooling 167
- storage system
 - architecture 5
 - implementation 5
 - service clearances 145
- storage-enclosure fillers 88
- storage, consolidating 50
- System i
 - control switch settings 175

T

- T10 DIF
 - ANSI support 51
 - Data Integrity Field 51
 - FB LUN 51
 - FB LUN protection 51
 - FB volume 51
 - Linux on z Systems 51
 - SCSI end-to-end 51
 - standard protection 51
 - z Systems 51
- Taiwan BSMI certificate 113, 114
- terminology v
- thermal load 162
- thin provisioning 120
 - ESE capacity controls 60
 - overview 59
- thin provisioning and Easy Tier 61
- three tiers 61
- tiles, perforated for cooling 167
- Tivoli Storage Productivity Center 46
- top exit
 - bracket 133, 135
 - measurements 133, 135
 - overhead cable management 133, 135
 - placement 133, 135
- top exit bracket
 - feature codes 134
- trademarks 212
- transparent cloud tiering 57, 122
- Transparent cloud tiering
 - adapters feature codes 102, 207

U

- understanding fixed block (FB)
 - architecture 51
- understanding logical volumes 52
- user interfaces 35

V

- VMware
 - Array Integration support 64
 - restrictions 64
- volume capacity
 - overview 53
- volume deletion
 - force option 59
 - safe option 59
- volume migration 61

- volume rebalancing 61
- volumes
 - allocation 34, 52
 - deletion 52
 - force option 52
 - modification 52

W

- warranty 195
- websites v
- weight
 - floor load capacity 138
 - reducing shipment 114
 - feature code 114
 - storage system, installed 142
- weight and dimensions
 - shipping container 128
- worksheets
 - provided equipment 197
 - customization 209
- WWID, switch settings 175

X

- XRC (Extended Remote Copy) (see z/OS Global Mirror) 125

Z

- z-synergy Services license 121
 - High Performance FICON 122
 - HyperPAV 122
 - Parallel Access Volumes 122
 - transparent cloud tiering 122
- z/OS Distributed Data Backup 122
- zHyperLink 123
- z/OS Global Mirror 125
- z/OS Metro/Global Mirror Incremental Resync 125
- zHyperLink 101
 - for previous models 206
 - I/O adapters feature codes 101
 - z-synergy feature 123
- zHyperLink cable
 - feature codes 102



Printed in USA

GC27-8525-16

