

IBM DS8882F Rack Mounted storage system
Version 8 Release 5.3

Introduction and Planning Guide



Note

Before using this information and the product it supports, read the information in [“Safety and environmental notices”](#) on page 116 and [“Notices”](#) on page 115.

This edition applies to version 8, release 5, modification 3 of the DS8882F Rack Mounted storage system.

This edition replaces GC27-9259-01.

© **Copyright International Business Machines Corporation 2018, 2019.**

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

Contents

About this book.....	vii
Who should use this book.....	vii
Conventions and terminology.....	vii
Publications and related information.....	vii
IBM Publications Center.....	xi
Feedback.....	xii
 Summary of changes	 xiii
 Chapter 1. Overview.....	 1
Machine types.....	2
Hardware.....	3
DS8882F.....	3
High Performance Flash Enclosures Gen2 pair.....	5
I/O enclosure.....	5
Processor nodes.....	6
Management enclosure.....	6
Power.....	6
Functions.....	7
Logical configuration.....	9
Logical configuration with DS8000 Storage Management GUI.....	9
Logical configuration with DS CLI.....	11
RAID implementation.....	13
Logical subsystems.....	14
Allocation methods.....	14
Management interfaces.....	15
DS8000 Storage Management GUI.....	15
DS command-line interface	16
DS Open Application Programming Interface	16
RESTful API.....	17
IBM Spectrum Control.....	17
IBM Copy Services Manager.....	17
Supported web browsers.....	18
 Chapter 2. Hardware features.....	 21
Storage complexes.....	23
Management console.....	24
Hardware specifics.....	24
Storage system structure.....	24
Flash drives.....	24
Drive maintenance policy.....	25
Host attachment overview.....	25
I/O load balancing.....	26
Storage consolidation.....	27
Count key data.....	27
Fixed block.....	27
T10 DIF support.....	27
Logical volumes.....	28
Allocation, deletion, and modification of volumes.....	28
LUN calculation.....	29

Extended address volumes for CKD.....	30
Quick initialization.....	30
Chapter 3. Data management features.....	31
Transparent cloud tiering.....	31
Dynamic volume expansion.....	32
Count key data and fixed block volume deletion prevention	33
Thin provisioning.....	33
ESE capacity controls for thin provisioning.....	33
IBM Easy Tier.....	34
VMware vStorage API for Array Integration support	35
Performance for IBM Z	36
Copy Services.....	37
Disaster recovery through Copy Services.....	44
Resource groups for Copy Services scope limiting.....	45
Comparison of Copy Services features.....	46
I/O Priority Manager.....	47
Securing data.....	48
Chapter 4. Storage system physical configuration.....	49
Configuration controls.....	49
Determining physical configuration features.....	49
Selecting management console features.....	50
Primary and secondary management consoles.....	50
Selecting storage features.....	50
Drive sets.....	51
Storage-enclosure fillers.....	51
Storage feature configuration rules.....	52
Physical and effective capacity.....	52
Selecting I/O adapter features.....	54
I/O enclosure.....	54
PCIe adapters.....	55
Fibre Channel host adapters and cables.....	55
Host adapter configuration rules.....	56
Selecting processor complex features.....	57
Transparent cloud tiering adapters.....	57
Processor licenses.....	57
Selecting processor memory features.....	57
System memory.....	57
Selecting power features.....	58
Power cords.....	58
Input voltage.....	61
Power feature configuration rules.....	61
Selecting other configuration features.....	61
BSMI certificate (Taiwan).....	62
Non-encryption certification key (China and Russia).....	62
Chapter 5. Planning use of licensed functions.....	63
Licensed function indicators.....	63
License scope.....	63
Ordering licensed functions.....	64
Rules for ordering licensed functions.....	65
Base Function license.....	66
Database protection.....	66
Encryption Authorization.....	67
IBM Easy Tier.....	67
I/O Priority Manager.....	67

Operating environment license.....	67
Thin provisioning	67
z-synergy Services license.....	67
High Performance FICON for z Systems.....	68
IBM HyperPAV	68
Parallel Access Volumes	68
Transparent cloud tiering.....	69
z/OS Distributed Data Backup.....	69
Copy Services license.....	69
Remote mirror and copy functions.....	70
FlashCopy function (point-in-time copy).....	70
Safeguarded Copy.....	70
z/OS Global Mirror.....	71
z/OS Metro/Global Mirror Incremental Resync.....	71
Copy Services Manager on the Hardware Management Console license.....	71
Chapter 6. Delivery and installation requirements.....	73
Acclimation.....	73
Shipment weights and dimensions.....	73
Receiving delivery.....	74
Installation site requirements.....	74
Planning the rack configuration.....	74
Planning for floor and space requirements.....	76
Planning for power requirements.....	77
Planning for environmental requirements.....	79
Planning for safety.....	82
Planning for network and communications requirements.....	83
Chapter 7. Storage complex setup.....	87
Company information.....	87
Management console network settings.....	87
Remote support settings.....	88
Notification settings.....	88
Power control settings.....	89
Control switch settings.....	89
Chapter 8. Data migration.....	91
Selecting a data migration method.....	92
Chapter 9. Security.....	95
Encryption.....	95
Encryption-key servers.....	95
Key lifecycle managers.....	96
Planning for full-disk encryption activation.....	96
User accounts and passwords.....	97
User accounts.....	97
Service accounts.....	97
NIST SP 800-131A security conformance.....	98
Chapter 10. License activation and management.....	99
Planning your licensed functions.....	99
Activation of licensed functions.....	100
Activating licensed functions.....	100
Scenarios for managing licensing.....	101
Adding storage to your machine.....	101
Managing a licensed feature.....	102

Appendix A. Accessibility features	103
Appendix B. Warranty information.....	105
Appendix C. IBM equipment and documents	107
Installation components.....	107
Customer components.....	108
Service components.....	108
Appendix D. Customization worksheets.....	109
Appendix E. Compliance standards.....	111
Appendix F. European Union General Data Protection Regulation Considerations.	113
Notices.....	115
Trademarks.....	116
Homologation statement.....	116
Safety and environmental notices.....	116
Safety notices and labels.....	116
Vendor-specific uninterruptible power supply safety statements.....	126
Environmental notices.....	127
Electromagnetic compatibility notices.....	127
Index.....	131

About this book

This book describes how to plan for a new installation of the DS8882F Rack Mounted storage system. It includes information about planning considerations, customization guidance, and configuration.

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 \(get.adobe.com/reader/\)](http://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.	V8.5.3 GC27-9259-02 V8.5.1 GC27-9259-01 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.	V8.5.3 GC27-8525-18 V8.5.1 GC27-8525-17 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.	V8.5.1 SC27-8526-10 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
Warranty Information for DS8000 series	http://www.ibm.com/support/docview.wss?uid=ssg1S7005239
IBM Safety Notices	Search for G229-9054 on the IBM Publications Center website
IBM Systems Environmental Notices	http://ibm.co/1fBgWFI
International Agreement for Acquisition of Software Maintenance (Not all software will offer Software Maintenance under this agreement.)	http://ibm.co/1fBmKPz
License Agreement for Machine Code	http://ibm.co/1mNiW1U
Other Internal Licensed Code	http://ibm.co/1kvABXE

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 https://www.ibm.com/support/docview.wss?uid=isg3T1025361	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 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.
DS8800 Code Bundle Information website(www.ibm.com/support/docview.wss?uid=s5g1S1003740)	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 <code>ver</code> command when you specify the <code>-l</code> parameter.
DS8870 Code Bundle Information website (www.ibm.com/support/docview.wss?uid=s5g1S1004204)	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 <code>ver</code> command when you specify the <code>-l</code> parameter.
DS8880 Code Bundle Information website (www.ibm.com/support/docview.wss?uid=s5g1S1005392)	Find information about code bundles for DS8880. The version of the currently active installed code bundle displays with the DS CLI <code>ver</code> command when you specify the <code>-l</code> 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 \(ibm.com/shop/publications/order\)](http://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, Modification 3 introduces the following new features.

Version 8.5.3

This table provides the current technical changes and enhancement to the IBM DS8000 as of June 7, 2019. Changed and new information is indicated by a vertical bar (|) to the left of the change.

Function	Description
Support for 1.92 TB Flash Tier 2 drives	See “Feature codes for drive sets” on page 51 for more information.

Chapter 1. Overview

The IBM DS8882F Rack Mounted storage system is a high-performance storage system that supports continuous operation, data security, and data resiliency. For high-availability, the hardware components are redundant.

DS8882F adds a modular rack-mountable enterprise storage system to the 533x all-flash machine type family. The modular system can be integrated into 16U contiguous space of an existing IBM z14™ Model ZR1 (z14 Model ZR1), IBM LinuxONE Rockhopper II (z14 Model LR1), or other standard 19-inch wide rack that conforms to EIA 310D specifications. The DS8882F allows you to take advantage of the DS8880 advanced features while limiting datacenter footprint and power infrastructure requirements. The modular system contains processor nodes, an I/O Enclosure, High Performance Flash Enclosures Gen2, a Management Enclosure (which includes the HMCs, Ethernet Switches, and RPCs), and battery backup modules to power the DS8882F modules. The HMCs are small form factor computers.

The DS8884, DS8886, DS8884F, DS8886F, and DS8888F systems are not documented in this publication. For information on those systems, refer to the *DS8880 Introduction and Planning Guide* (GC27-8525-16).

Licensed functions are available in four groups:

Base Function

The Base Function license is required for each DS8882F 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 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

DS8882F 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 DS8882F. Order a hardware machine type for the storage system and a corresponding function authorization machine type for the licensed functions that are planned for use.

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

<i>Table 4. Available hardware and function-authorization machine types.</i>			
Hardware		Licensed functions	
Hardware machine type	Available hardware models	Corresponding function authorization machine type	Available function authorization models
5331 (1-year warranty period)	983	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)	

The machine types for the DS8882F specify the service warranty period. The warranty is used for service entitlement checking when notifications for service are called home. The model 983 reports 2107 as the machine type to attached host systems.

Hardware

The architecture of the IBM DS8882F Rack Mounted storage system is based on three major elements that provide function specialization and three tiers of processing power.




Figure 1. DS8882F architecture

The DS8882F 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

DS8882F (machine type 533x model 983)

The DS8882F is an entry-level, high-performance storage system that includes only High Performance Flash Enclosures Gen2.

DS8882F storage system features 6-core processors and supports one High Performance Flash Enclosure Gen2 pair with up to 48 Flash Tier 0, Flash Tier 1, or Flash Tier 2 drives. This modular rack-mountable enterprise storage system can be integrated into 16U contiguous space of an IBM z14 Model ZR1 (z14 Model ZR1), IBM LinuxONE Rockhopper II (z14 Model LR1), or other standard 19-inch wide rack that conforms to EIA 310D specifications to take advantage of the DS8880 advanced features while limiting

datacenter footprint and power infrastructure requirements. The modular system contains processor nodes, an I/O Enclosure, High Performance Flash Enclosures Gen2, a Management Enclosure (which includes the HMCs, Ethernet Switches, and RPCs), and battery backup modules to power the DS8882F modules. The HMCs are small form factor computers.

Note: The standard 19-inch wide rack installation (feature code 0939) supports an optional 1U keyboard and display (feature code 1765). The 16U contiguous space requirement does not include space for the optional keyboard and display, but they are not required to reside contiguously with the DS8882F model 983. For accessibility, the keyboard and display must be mounted at a height of 15 - 46 inches. If you add the keyboard and display, ensure that you provide adequate space to accommodate them. The optional 1U keyboard and display are not available with the z14 model ZR1 installation (feature code 0937) or the z14 model LR1 installation (feature code 0938). Feature code 1765 requires one IEC C13 power connection. The model 983 power supplies have no available power output connectors. You must make one available on the C13 power connection in the rack.

The DS8882F uses 16 Gbps Fibre Channel host adapters that run Fibre Channel Protocol (FCP) or FICON protocol. The High Performance FICON (HPF) feature is also supported.

The DS8882F 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 DS8882F.

Table 5. Components for the DS8882F (model 983)					
Processors	System memory	Processor memory	I/O enclosure	Host adapters (4 port)	High Performance Flash Enclosure Gen2 pairs ¹
6-core	64 GB	32 GB	1	2 or 4	1
6-core	128 GB	64 GB	1	2 or 4	1
6-core	256 GB	128 GB	1	2 or 4	1
1. High Performance Flash Enclosures Gen2 are installed in pairs.					

Table 6. Maximum capacity for the DS8882F (model 983)			
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 drives
6-core	64 GB	48	737 TB
6-core	128 GB	48	737 TB
6-core	256 GB	48	737 TB

DS8882F (model 983) overview

The DS8882F (model 983) consists of eight 2U modules for installation in an existing rack.

The model 983 includes the following components:

- High Performance Flash Enclosure Gen2 pair
- I/O enclosure
- Two processor nodes (available with POWER8 processors)
- Management enclosure
- Two 3 kVA battery backup modules




Figure 2. DS8882F model 983 modules

High Performance Flash Enclosures Gen2 pair

The High Performance Flash Enclosure Gen2 is a 2U storage enclosure that is installed in pairs. For version 8.5, the model 983 supports only one High Performance Flash Enclosure Gen2 pair.

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
 - 24 2.5-inch flash drives (or drive fillers)

I/O enclosure

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.

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 contains the following adapters:

Flash RAID adapters

PCIe-attached adapter with four SAS ports. The Flash RAID adapters (imbedded in the PCIe adapter) provide RAID controllers and SAS connectivity to the High Performance Flash Enclosures Gen2.

Host adapters

An I/O enclosure can support two or four 4-port host adapters, providing eight or 16 host ports.

Each of the four 16 Gbps Fibre Channel ports on a PCIe-attached host adapter 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.

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.

Management enclosure

The model 983 contains a management enclosure.

The management enclosure contains the following components:

- Two Hardware Management Consoles (HMCs)
- Two Ethernet switches
- Two power control cards
- Two power supply units (PSUs) to power the management enclosure
- One Local/Remote switch assembly

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 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.

Ethernet switches

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

Power

Two redundant 3 kVA battery backup modules supply 230 V AC power to the DS8882F storage system. Each battery backup module receives input power from a single-phase line cord.

If both battery backup modules lose input power, they have sufficient capacity to continue to supply AC power to the DS8882F until it has completed a fire hose dump to protect modified data. The system will then gracefully power off.

Functional overview

The following list provides an overview of some of the features that are associated with DS8882F.

Note: Some storage system functions are not available or are not supported in all environments. See the [IBM System Storage Interoperation Center \(SSIC\) website \(www.ibm.com/systems/support/storage/config/ssic\)](http://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

DS8882F supports full redundancy, but some components are a single point of repair. It is designed to support nondisruptive changes: repair, and licensed function upgrades. In addition, logical configuration changes can be made nondisruptively. For example:

- 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.

Energy reporting

You can use the DS8882F 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 **shows** command. For more information, see the **shows** 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 DS8882F 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 DS8882F system.

Setting up unified sign-on support for the DS8882F 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 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 DS8882F system to automatically balance I/O access to drives to avoid hot spots on arrays.

Easy Tier 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 DS8882F 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

Copy Services

The DS8882F 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
 - Multiple 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 volume. The I/O Priority Manager detects when a higher-priority application is hindered by a lower-priority application that is competing for the same system resources. This detection might occur when multiple applications request data from the same drives. When I/O Priority Manager encounters this situation, it delays lower-priority I/O data to assist the more critical I/O data in meeting its performance targets.

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

The default setting for this feature is disabled.

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

Restriction of hazardous substances (RoHS)

The DS8882F 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 3 on page 10](#) illustrates the concepts of configuration.




Figure 3. 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 4 on page 12](#) 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 4. 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 1.92 TB Flash Tier 2 and 3.84 TB Flash Tier 1 drives. RAID 6 is the only supported RAID type for 7.68 TB and 15.36 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 1.92 TB Flash Tier 2 and 3.84 TB Flash Tier 1 drives. RAID 6 is the only supported RAID type for 7.68 TB and 15.36 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 1.92 TB Flash Tier 2 and 3.84 TB Flash Tier 1 drives. RAID 6 is the only supported RAID type for 7.68 TB and 15.36 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 1.92 TB Flash Tier 2 and 3.84 TB Flash Tier 1 drives. RAID 6 is the only supported RAID type for 7.68 TB and 15.36 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 Spectrum Control and IBM Copy Services Manager.

DS8000 Storage Management GUI

Use the DS8000 Storage Management GUI to configure and manage storage, and to 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 18.

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 your storage system 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 development, testing, and debugging of the client management infrastructures in your storage system. 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 IBM storage management software
- Providing greater synergy between storage management software and IBM storage devices
- Reducing the number of servers that are required to manage your software infrastructure
- Migrating from basic device management to storage management applications that provide higher-level functions

For more information, see [IBM Spectrum Control online product documentation in IBM Knowledge Center \(www.ibm.com/support/knowledgecenter\)](http://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 7. 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 the IBM DS8000 series online product documentation for security best practices.

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 8. Feature codes for hardware features		
Feature code	Feature	Description
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.
0937	zFlex Frame field merge	Indicates that the DS8882F will be installed in an existing IBM z14 model ZR1 rack
0938	Rockhopper II field merge	Indicates that the DS8882F will be installed in an existing IBM z14 model LR1 rack
0939	Customer Rack field merge	Indicates that the DS8882F will be installed in an existing standard 19-inch wide rack that conforms to EIA 310D specifications
1021	Single-phase power cord, 250 V, 20 A	NEMA L6-20P
1022	Single-phase power cord, 250 V, 16 A	CEE 7 VII
1023	Single-phase power cord, 250 V, 16 A	SANS 164
1024	Single-phase power cord, 250 V, 16 A	CEI 23-16
1025	Single-phase power cord, 250 V, 20 A	RS 3720DP
1026	Single-phase power cord, 250 V, 16 A	IEC 309
1027	Single-phase power cord, 250 V, 15 A	AS/NZS 3112
1028	Single-phase power cord, 250 V, 15 A	JIS C 8303 6-20P
1029	Single-phase power cord, 125 - 250 V, 16 A	IEC 60320-2-2
1030	Single-phase power cord, 250 V, 20 A	IRAM 2073
1031	Single-phase power cord, 250 V, 16 A	KSC 8305
1032	Single-phase power cord, 250 V, 16 A	IS 6538
1033	Single-phase power cord, 250 V, 16 A	GB 2099.1, 1002
1034	Single-phase power cord, 250 V, 20 A	NBR 14136
1035	Single-phase power cord, 250 V, 20 A	CNS 10917-3
1036	Single-phase power cord, 250 V, 16 A	SI 32
1037	Single-phase power cord, 250 V, 16 A	SEV 1011

Table 8. Feature codes for hardware features (continued)

Feature code	Feature	Description
1057	Battery backup module (one)	Two battery backup modules are required
1101	5 ft. ladder	For servicing components at a height of 35U or more
1102	3 ft. platform ladder	For servicing components at a height of 20U or more
1103	Rolling step stool	For all models
1151	Secondary management console	Redundant management console for high availability For model 983, this feature is required
1241	Drive enclosure pair total	Admin feature totaling all disk enclosure pairs installed in the model
1305	2U I/O enclosure pair	
1410	Fibre Channel cable	40 m (131 ft), 50 micron OM3 or higher, multimode LC
1413	Fibre Channel cable	3 m (10 ft), 50 micron OM3 or higher, multimode LC
1420	Fibre Channel cable	31 m (102 ft), 9 micron OS1 or higher, single mode LC
1423	Fibre Channel cable	3 m (10 ft), 9 micron OS1 or higher, single mode LC
1603	High Performance Flash Enclosure Gen2 pair	For flash drives
1610	400 GB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives) No intermix with Flash Tier 1 or Flash Tier 2 drive sets
1611	800 GB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives) No intermix with Flash Tier 1 or Flash Tier 2 drive sets
1612	1.6 TB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives) No intermix with Flash Tier 1 or Flash Tier 2 drive sets
1613	3.2 TB 2.5-inch Flash Tier 0 drive set	Flash drive set (16 drives) No intermix with Flash Tier 1 or Flash Tier 2 drive sets
1622	1.92 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

Table 8. Feature codes for hardware features (continued)

Feature code	Feature	Description
1623	3.84 TB 2.5-inch Flash Tier 1 drive set	Flash drive set (16 drives) No intermix with Flash Tier 0 drive sets
1624	7.68 TB 2.5-inch Flash Tier 2 drive set	Flash drive set (16 drives) No intermix with Flash Tier 0 drive sets
1625	15.36 TB 2.5-inch Flash Tier 2 drive set	Flash drive set (16 drives) No intermix with Flash Tier 0 drive sets
1699	High Performance Flash Enclosure Gen2 filler set	Includes 16 fillers
1765	Optional 1U keyboard and display	Not available with the z14 model ZR1 installation (feature code 0937) or the z14 model LR1 installation (feature code 0938)
1885	DS8000 Licensed Machine Code R8.5	Microcode bundle 88.x.xx.x for base model 983
3065	Base I/O expander/Flash RAID adapter (required)	Required to support one High Performance Flash Enclosure Gen2 pair and the first pair of host adapters
3066	I/O expander (optional)	Required to support the second pair of host adapters in the I/O Enclosure
3354	Fibre Channel host-adapter	4-port, 16 Gbps shortwave FCP and FICON host adapter PCIe Minimum of two can be ordered Requires feature code 3065 or 3066
3454	Fibre Channel host-adapter	4-port, 16 Gbps longwave FCP and FICON host adapter PCIe Minimum of two can be ordered
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
4233	64 GB system memory	(6-core)
4234	128 GB system memory	(6-core)
4235	256 GB system memory	(6-core)
4421	6-core POWER8 processors	Requires feature code 4233, 4234, or 4235

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. Model 983 provides a second management console.

Management console

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

The *management console* is a dedicated processor unit that 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 (already provided with model 983).

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 pair of storage enclosures.

Storage server

The storage server consists of two processor complexes 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 (for model 983, cards are located in the Management enclosure) coordinate the power management within the storage unit. The power control cards are attached to the service processors in each processor complex.

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
 - 3.84 TB
- 2.5-inch Flash Tier 2 drives with FDE
 - 1.92 TB
 - 7.68 TB

– 15.36 TB

Note: Intermix of high performance drives (Flash Tier 0) with high capacity drives (Flash Tier 1 or Flash Tier 2) is not supported within an enclosure pair.

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

The DS8882F model 983 provides only 4-port 16 Gbps host adapters, and a maximum of 16 ports is available.

DS8882F model 983

The following table shows the host adapter plug order.

Table 9. Plug order for 4-port HA slots for the I/O enclosure .				
Host adapter pair	Slot number			
	C3	C4	C5	C6
First host adapter pair	1			1
Second host adapter pair		2	2	

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. 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. 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 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 \(www.ibm.com/systems/support/storage/config/ssic\)](http://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.

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.

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 10 on page 29 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 10. 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 \(www.ibm.com/support/knowledgecenter/SSWPVP/\)](http://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. With Safeguarded Copy, *in use* means that the volumes have data saved in the backup repository.

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. However, an *in use* volume that has Safeguarded Copy space cannot be deleted with the **-force** option.

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 homogeneous pools, hot spot management, rank depopulation, manual volume migration, and thin provisioning support (ESE volumes only).

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.

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.

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.

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, reprotection, and failback. The supported Copy Services are Metro Mirror, Global Mirror, Metro-Global Mirror, and FlashCopy.

The IBM Storage Management Console plug-in enables VMware administrators to manage their systems from within the VMware management environment. This plug-in provides an integrated view of IBM storage to VMware virtualize datastores that are required by VMware administrators. For information, see the IBM Storage Management Console for VMware vCenter (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

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 69](#).

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.

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 from a DS8882F 4-port host adapter to a previous release DS8000 (Release 6.0 or later), which supports 8-port host adapters, you can only connect the lower four ports of the 8-port host adapter.
- 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 5 on page 40 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 5. 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 \(www.ibm.com/systems/support/storage/config/ssic\)](http://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 6 on page 46 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 6. 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 11 on page 46 provides a brief summary of the characteristics of the Copy Services features that are available for the storage system.

Table 11. 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

<i>Table 11. Comparison of features (continued)</i>			
Feature	Description	Advantages	Considerations
Metro/Global Mirror	Three-site, long distance disaster recovery replication	A backup site is maintained regardless of which one of the sites is lost.	Recovery point objective (RPO) might grow if bandwidth capability is exceeded.
Metro Mirror	Synchronous data copy at a distance	No data loss, rapid recovery time for distances up to 300 km.	Slight performance impact.
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.

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 12 on page 47](#) lists performance groups that are predefined and have the associated performance policies:

<i>Table 12. Performance groups and policies.</i>		
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

Table 12. Performance groups and policies. (continued)

Performance group ¹	Performance policy	Performance policy description
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. The DS8880 systems use AES-256 encryption.

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

The IBM DS8882F Rack Mounted storage system 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 has a valid configuration. There is no charge for these features.

Your storage system can include the following indicators:

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 models of which your storage system is to be comprised.
3. Order a primary and secondary management console for each storage system.
4. For each storage system, 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.
5. Determine the I/O adapter features that you need for your storage system.
6. Determine the appropriate processor memory feature code that is needed.
7. Decide which power features that you must order.
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.

The primary and secondary management console are included in the DS8882F model 983.

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 consoles (also known as the Hardware Management Consoles or HMCs) are dedicated appliances physically located inside the Management enclosure. 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.

The secondary HMC is a redundant management console for environments with high-availability requirements, and is required for model 983.

Feature codes for management consoles

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

The primary management console is included by default in DS8882F storage systems.

Table 13. Feature codes for management consoles		
Feature code	Description	Models
1151	Secondary management console	Redundant management console for high availability For model 983, this feature is required

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
- Enclosure filler features

Feature codes for drive sets

Use these feature codes to order sets of encryption capable flash drives.

The flash drives can be installed only in High Performance Flash Enclosures Gen2. The High Performance Flash Enclosure Gen2 pair can contain 16, 32, or 48 flash drives. All flash drives in a High Performance Flash Enclosure Gen2 pair must be the same type.

Table 14. 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 capable 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, 2}
1613	3.2 TB	2.5-in. Flash Tier 0 drives	16	N/A	Yes	6, 10 ^{1, 2}
1622	1.92 TB	2.5-in. Flash Tier 2 drives	16	N/A	Yes	6, 10 ²
1623	3.84 TB	2.5-in. Flash Tier 1 drives	16	N/A	Yes	6, 10 ^{1, 2}
1624	7.68 TB	2.5-in. Flash Tier 2 drives	16	N/A	Yes	6 ^{1, 2}
1625	15.36 TB	2.5-in. Flash Tier 2 drives	16	N/A	Yes	6 ^{1, 2}

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 sales representative.
2. RAID 6 is the default RAID type for all drives larger than 1 TB, and it is the only supported RAID type for 7.68 TB drives and 15.36 TB drives.
3. Within a High Performance Flash Enclosure Gen2 pair, no intermix of high performance drives (Flash Tier 0) with high capacity drives (Flash Tier 1 or Flash Tier 2) is supported.

Storage-enclosure fillers

Storage-enclosure fillers fill empty drive slots in the storage enclosures. The fillers ensure sufficient airflow across populated storage.

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 High Performance Flash Enclosures Gen2.

Table 15. Feature codes for storage enclosures	
Feature code	Description
1699	Filler set for 2.5-in. High Performance Flash Enclosures Gen2; includes 16 fillers

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.

Flash drive sets

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

Storage enclosure fillers

For the 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. If two drive sets are ordered (32 drives), one filler set is required to fill the remaining 16 slots. Each drive slot in a High Performance Flash Enclosures Gen2 must have either a flash drive or a filler.

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 High Performance Flash Enclosures Gen2, there are 16 identical flash drives per drive set, up to three drive sets per enclosure pair.

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 DS8882F 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, 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.

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 DS8882F

Use the following information to calculate the physical and effective capacity for High Performance Flash Enclosures Gen2.

The default RAID type for all drives larger than 1 TB is RAID 6, and it is the only RAID type supported for 7.68 TB and 15.36 TB drives. 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 16. RAID capacities for High Performance Flash Enclosures Gen2.								
Flash drive disk size	Physical capacity of Flash drive 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 16. RAID capacities for High Performance Flash Enclosures Gen2. (continued)								
Flash drive disk size	Physical capacity of Flash drive 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
1.92 TB	15.4 TB	FB Lg Ext	5168	6902	n/a	n/a	8636	10370
		FB Sm Ext	330783	441769	n/a	n/a	552748	663727
		CKD Lg Ext	5796	7741	n/a	n/a	9686	11631
		CKD Sm Ext	307231	410315	n/a	n/a	513392	616474
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
3.84 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
7.68 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
15.36 TB	246 TB	FB Lg Ext	n/a	n/a	n/a	n/a	68980	82782
		FB Sm Ext	n/a	n/a	n/a	n/a	4414735	5298103
		CKD Lg Ext	n/a	n/a	n/a	n/a	77365	92846
		CKD Sm Ext	n/a	n/a	n/a	n/a	4100392	4920882

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:

- 2U I/O enclosure
- Flash RAID adapter pair
- Host adapters
- Host adapters Fibre Channel cables

I/O enclosure

The I/O enclosure holds the I/O adapters and provides connectivity between the I/O adapters and the storage processors.

The I/O adapters in the I/O enclosures can be either Flash RAID adapters or host adapters. Each I/O enclosure can support up to two Flash RAID adapters (one pair), which are imbedded into the PCIe adapter, and two or four host adapters installed in pairs (not to exceed 16 host adapter ports).

Feature code for 2U I/O enclosure

Use this feature code to identify the 2U I/O enclosure for your storage system.

This I/O enclosure feature includes one 2U I/O enclosure. This feature supports two or four 16 Gbps 4-port host adapters (installed in pairs).

Table 17. Feature codes for I/O enclosures	
Feature code	Description
1305	2U I/O enclosure

Feature codes for PCIe adapters

Use these feature codes to identify Flash RAID adapters for your storage system.

Table 18. Feature codes for PCIe adapters		
Feature code	Feature name	Description
3065	Base I/O expander/ Flash RAID adapter (required)	Required to support one High Performance Flash Enclosure Gen2 pair and the first pair of host adapters
3066	I/O expander (optional)	Required to support the second pair of host adapters in the I/O Enclosure

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. Fibre Channel host adapters are installed in the I/O enclosure.

Adapters are 4-port 16 Gbps.

Supported protocols include the following types:

- SCSI-FCP upper layer protocol (ULP) on point-to-point and fabric.

Note: The 16 Gbps adapter does not support arbitrated loop topology at any speed.

- FICON ULP on point-to-point and fabric topologies.

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.

Table 19. Feature codes for Fibre Channel host adapters		
Feature code	Description	Receptacle type
3354	4-port, 16 Gbps shortwave FCP and FICON host adapter, PCIe	LC
3454	4-port, 16 Gbps longwave FCP and FICON host adapter, PCIe	LC

Feature codes for Fibre Channel cables

Use these feature codes to order Fibre Channel cables to connect Fibre Channel host adapters. Take note of the distance capabilities for cable types.

Table 20. 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 adapter (feature code 3354)
1413	50 micron OM3 or higher Fibre Channel cable, multimode	3 m (10 ft)	
1420	9 micron OS1 or higher Fibre Channel cable, single mode	31 m (102 ft)	• Longwave Fibre Channel or FICON adapter (feature code 3454)
1423	9 micron OS1 or higher Fibre Channel cable, single mode	3 m (10 ft)	

Table 21. Multimode cabling limits.

Fibre cable type	Distance limits relative to 16 Gbps
OM1 (62.5 micron)	Not recommended
OM2 (50 micron)	35 m, but not recommended
OM3 (50 micron)	100 m
OM4 (50 micron)	125 m

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:

- How many host adapters will I install? You can install either two or four in fixed locations in pairs.
- 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 guideline.

- Longwave and shortwave host adapter intermix is only supported if a pair of each is ordered.

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. The DS8882F contains two processor enclosures (POWER8 servers) that contain the processors and memory that drives all functions in the storage system.

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 22. 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 983

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 23. Feature codes for processor licenses

Feature code	Description	Corequisite feature code for memory
4421	6-core POWER8 processor feature	4233, 4234, or 4235

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 the DS8882F.

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 24. Feature codes for system memory.

Feature code	Description	Model
4233 ¹	64 GB system memory	983 (6-core)
4234 ¹	128 GB system memory	983 (6-core)
4235 ¹	256 GB system memory	983 (6-core)

Table 24. Feature codes for system memory. (continued)		
Feature code	Description	Model
Notes: 1. Feature codes 4233, 4234, and 4235 require 6-core processor license feature code 4421.		

Power features

You must specify the power features to include on your storage system.

Power cords

A pair of power cords (also known as *power cables*) is required.

Feature codes for power cords

Use these feature codes to order power cords. Ensure that you meet the requirements for each power cord and connector type that you order.

For the z14 model ZR1 installation (feature code 0937) or the z14 model LR1 installation (feature code 0938), the DS8882F model 983 power cords are connected to facility outlets outside the ZR1 or LR1 rack.


Table 25. Feature codes for power cords		
Feature code	Power cord facility plug type	Plug diagram
1021	Single-phase 4.3 m (14 ft) power cord, 250 V, 20 A NEMA L6-20P	
1022	Single-phase 4.3 m (14 ft) power cord, 250 V, 16 A CEE 7 VII	
1023	Single-phase 4.3 m (14 ft) power cord, 250 V, 16 A SANS 164	
1024	Single-phase 4.3 m (14 ft) power cord, 200 - 240 V, 16A CEI 23-16	
1025	Single-phase 4.3 m (14 ft) power cord, 250 V, 20 A RS 3720DP	

Table 25. Feature codes for power cords (continued)







Feature code	Power cord facility plug type	Plug diagram
1026	Single-phase 4.3 m (14 ft) power cord, 250 V, 16 A IEC 309	
1027	Single-phase 4.3 m (14 ft) power cord, 250 V, 15 A AS/NZS 3112	
1028	Single-phase 4.3 m (14 ft) power cord, 250 V, 15 A JIS C 8303 6-20P	
1029	Single-phase 2.5 m (8 ft) power cord, 125 - 250 V, 16 A IEC 60320-2-2 Note: This is the default power cord. Typically, it is used to attach to rack PDU with C20 outlets. Before ordering, verify the power cord is compatible with available rack or facility power outlets.	
1030	Single-phase 4.3 m (14 ft) power cord, 250 V, 20 A IRAM 2073	
1031	Single-phase 4.3 m (14 ft) power cord, 250 V, 16 A KSC 8305	
1032	Single-phase 4.3 m (14 ft) power cord, 250 V, 16 A IS 6538	
1033	Single-phase 4.3 m (14 ft) power cord, 250 V, 16 A GB 2099.1, 1002	
1034	Single-phase 4.3 m (14 ft) power cord, 250 V, 20 A NBR 14136	

Table 25. Feature codes for power cords (continued)		
Feature code	Power cord facility plug type	Plug diagram
1035	Single-phase 4.3 m (14 ft) power cord, 250 V, 20 A CNS 10917-3	
1036	Single-phase 4.3 m (14 ft) power cord, 250 V, 16 A SI 32	
1037	Single-phase 4.3 m (14 ft) power cord, 250 V, 16 A SEV 1011	

The following list provides standard plug types and the countries in which they are commonly used. You can use a plug standard that is not identified here as common to your country. For example, NEMA L6-20P, RS 3720DP, or IEC309 locking plugs might be preferred and can be used in most countries.

NEMA L6-20P

United States, Canada

CEE 7 VII

Afghanistan, Albania, Algeria, Andorra, Angola, Armenia, Austria, Azerbaijan, Belarus, Belgium, Benin, Bosnia and Herzegovina, Bulgaria, Burkina Faso, Burundi, Cambodia, Cameroon, Cape Verde, Central African Republic, Chad, Comoros, Congo (Democratic Republic of), Congo (Republic of), Cote D'Ivoire (Ivory Coast), Croatia (Republic of), Czech Rep, Dahomey, Djibouti, Egypt, Equatorial Guinea, Eritrea, Estonia, Ethiopia, Macedonia (former Yugoslav Republic of), Madagascar, Mali, Martinique, Mauritania, Mauritius, Mayotte, Moldova (Republic of), Monaco, Mongolia, Morocco, Mozambique, Netherlands, New Caledonia, Niger, Norway, Poland, Portugal, Reunion, Romania, Russian Federation, Rwanda, Sao Tome and Principe, Saudi Arabia, Senegal, Serbia, Slovakia, Slovenia (Republic of), Somalia, Spain, Suriname, Sweden, Syrian Arab Replublic, Tajikistan, Tahiti, Togo, Tunisia, Turkey, Turkmenistan, Ukraine, Upper Volta, Uzbekistan, Vanuatu, Vietnam, Wallis and Futuna, Yugoslavia (Federal Republic of), ZaireFinland, France, French Guyana, French Polynesia, Gabon, Georgia, Germany, Greece, Guadeloupe, Guinea, Guinea Bissau, Hungary, Iceland, Indonesia, Iran, Kazakhstan, Kyrgyzstan, Laos (Peoples Democratic Republic of), Latvia, Lebanon, Lithuania, Luxembourg

SANS 164

Bangladesh, Lesotho, Macao, Maldives, Namibia, Nepal, Pakistan, Samoa, South Africa, Sri Lanka, Swaziland

CEI 23-16

Chile, Italy, Libyan Arab Jamahiriya

RS 3720DP

United States, Canada

IEC 309

Denmark, Liechtenstein

AS/NZS 3112

Australia, Fiji, Kiribati, Nauru, New Zealand, Papua New Guinea

JIS C 8303 6-20P

Japan

IEC 60320-2-2

Worldwide

IRAM 2073

Argentina, Paraguay, Uruguay

KSC 8305

Korea (Democratic Peoples Republic of), Korea (Republic of)

IS 6538

India

GB 2099.1, 1002

China (SAR)

NBR 14136

Brazil

CNS 10917-3

Taiwan

SI 32

Israel

SEV 1011

Switzerland

Input voltage

The battery backup module distributes power that ranges from 200 V AC to 240 V AC.

Feature codes for battery backup modules

Use these feature codes to identify battery backup modules.

Table 26. Feature codes for battery backup modules		
Feature code	Description	Requirements
1057	Battery backup module	Two battery backup modules are included in the DS8882F model 983.

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 two power cord features.
- You must select the power cord that is appropriate for the input voltage and outlet type of the storage system in the frame that the model 983 is located.

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.

- BSMI certificate (Taiwan)
- Encryption not capable (China or Russia)

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 order the Bureau of Standards, Metrology, and Inspection (BSMI) certification documents that are required when the storage system is shipped to Taiwan.

Table 27. Feature code for the BSMI certification documents (Taiwan)	
Feature code	Description
0400	BSMI certification documents

Non-encryption certification key (China and Russia)

The encryption not capable feature for China and Russia disables the encryption capabilities of the storage system.

If the storage system that you order is shipped to China or Russia, you must order this option to ensure that the non-encryption certification key is applied to comply with government encryption requirements.

Feature code for non-encryption certification key (China and Russia)

Use this feature code to you order the non-encryption certification key that is required when the storage system is shipped to China or Russia.

Table 28. Feature code for non-encryption certification key (China and Russia)	
Feature code	Description
0403	Non-encryption certification key

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 for the licensed function before the feature activation code is provided. 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 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 \(www.ibm.com/storage/dsfa\)](http://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 29. Total physical capacity for drive-set features		
Drive sizes	Total physical capacity	Drives per feature
400 GB flash drives	6.4 TB	16
800 GB flash drives	12.8 TB	16
1.6 TB flash drives	25.6 TB	16
3.2 TB flash drives	51.2 TB	16
3.84 TB flash drives	61.4	16
7.68 TB flash drives	122 TB	16
15.36 TB flash drives	244 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 \(www.ibm.com/storage/dsfa\)](http://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 \(www.ibm.com/storage/dsfa\)](http://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 30. 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)

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 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 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

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 (capacity) that is authorized to use the product operating environment.

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

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 31. 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)

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.

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 32. 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)

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 includes the following feature codes.

<i>Table 33. Copy Services Manager on Hardware Management Console license feature codes</i>	
Feature Code	Feature code for licensed function indicator
8451	CSM on HMC - active

Chapter 6. Delivery and installation requirements

You must ensure that you properly plan for the delivery and installation of your DS8882F Rack Mounted 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 107.

Acclimation

Server and storage equipment must be gradually acclimated to the surrounding environment to prevent condensation.

When server and storage equipment is shipped in a climate where the outside temperature is below the dew point of 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 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.

“[Shipment weights and dimensions](#)” on page 73 shows the final packaged dimensions and maximum packaged weight of the storage system component shipments.

Table 34. Packaged dimensions and weight for storage systems (all countries).

Container	Packaged dimensions	Maximum packaged weight
DS8882F model 983	Height 1.49 m (58.7 in.) Width 1.05 m (41.3 in.) Depth 1.30 m (51.2 in.)	338 kg (745 lb)

Receiving delivery

The shipping carrier is responsible for delivering and unloading the storage system as close to its final destination as possible. You must ensure that your loading ramp and your receiving area can accommodate your storage system shipment.

About this task

Use the following steps to ensure that your receiving area and loading ramp can safely accommodate the delivery of your storage system:

Procedure

1. Find out the packaged weight and dimensions of the shipping containers in your shipment.
2. Ensure that your loading dock, receiving area, and elevators can safely support the packaged weight and dimensions of the shipping containers.

Installation site requirements

You must ensure that the location where you plan to install your storage system meets all requirements.

Planning the rack configuration

Ensure that the rack where you plan to install your storage system meets the rack requirements.

About this task

When you are planning the rack for your storage system, you must answer the following questions that relate to rack specifications and available space:

- Where are you installing the storage system? The DS8882F model 983 is a rack mountable system consisting of eight 2U modules. There are three different rack scenarios:
 - An existing IBM z14 Model ZR1 (z14 Model ZR1)
 - An existing IBM LinuxONE Rockhopper II (z14 Model LR1)
 - Other standard 19-inch wide rack that conforms to EIA 310D specifications:
 - 19-inch EIA rails
 - Minimum rail depth of 700 mm
 - Maximum rail depth of 780 mm
- Does the rack in which you are installing the DS8882F model 983 have adequate space to accommodate the components? The rack must have 16U contiguous space to mount the eight 2U modules.

Note: The standard 19-inch wide rack installation (feature code 0939) supports an optional 1U keyboard and display (feature code 1765). The 16U contiguous space requirement does not include space for the optional keyboard and display, but they are not required to reside contiguously with the DS8882F model 983. For accessibility, the keyboard and display must be mounted at a height of 15 - 46 inches. If you add the keyboard and display, ensure that you provide adequate space to accommodate them. The optional 1U keyboard and display are not available with the z14 model ZR1 installation (feature code 0937) or the z14 model LR1 installation (feature code 0938). Feature code 1765 requires one IEC C13 power connection. The model 983 power supplies have no available power output connectors. You must make one available on the C13 power connection in the rack.




Figure 7. DS8882F model 983 installation front view

Existing rack with
16U contiguous space

DS8882F model 983




Figure 8. DS8882F model 983 installation rear view

Planning for floor and space requirements

Ensure that the location where you plan to install your storage system meets space and floor requirements.

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:

- Does the floor of the installation site meet floor-load requirements for the existing z14 Model ZR1, z14 Model LR1, or other standard 19-inch wide frame, plus the additional weight of the DS8882F model 983?
- 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

Dimensions and weight

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 35. DS8882F model 983 dimensions and weights.		
Dimensions		Maximum weight
Height without optional keyboard and display	71.1 cm (28 in.)	242.7 kg (535 lb)
Height with optional keyboard and display	75.6 cm (29.8 in.)	
Distance between rack front and rear EIA rails	70.1 - 78 cm (27.6 - 30.7 in.)	
Minimum distance from cabinet rear rail to rear door	20.3 cm (8 in.)	
Minimum distance from cabinet front rail to front door	5 cm (2 in.)	

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.

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 the DS8882F model 983.

Important: To eliminate a single point of failure, independent power feeds to each DS8882F 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.

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 36. 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	14.4 Amps
Recommended wall breaker rating (1 ph)	15 - 20 Amps
Steady-state input frequencies	50 $\pm$ 3 or 60 $\pm$ 3.0 Hz
PLD input frequencies	< 20 milliseconds

The battery backup modules are each "Class I" rated for the following voltages and maximum currents for those voltages.

AC voltage 50 or 60 Hz	200 V	208 V	220 V	230 V	240 V
Maximum rated current	14.2 Amps	13.7 Amps	14.4 Amps	13.8 Amps	13.2 Amps

Power connector requirements

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

Find the power cord that is appropriate for your site. Ensure that your site meets the specified power connector requirements. For power cord feature codes, plug diagrams, and countries in which each plug type is commonly used, refer to ["Feature codes for power cords"](#) on page 58.



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 storage system. The storage system 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 15 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.

Power consumption and environmental information

When you are planning to meet the power requirements, 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.

<i>Table 37. Power consumption and environmental information for the DS8882F model 983.</i>		
Measurement	Unit of measure	Model 983
Peak electric power	kilovolt amperes (kVA)	3.0
Thermal load	British thermal units (BTU) per hour	10236
Capacity of exhaust	cubic meters per minute (cubic feet per minute or CFM)	22.7 (800)
Ground leakage current	milliamperes (mA)	11
Startup current	amperes (A or amp)	≤ 100
Startup current duration	microseconds (μs or μsec)	< 200

Acoustic declaration

<i>Table 38. Acoustic declaration for fully configured version 8.5 storage system.</i>				
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 983	7.9	7.9	62	61
Notes: 1. LWAd is the statistical upper-limit A-weighted sound power level (rounded to the nearest 0.1 B). 2. LpAm is the mean A-weighted emission sound pressure level that is measured at the 1-meter bystander positions (rounded to the nearest dB). 3. 10 dB (decibel) = 1 B (bel). 4. 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.

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 39 on page 79 provides the operating ranges for your storage system when its power is on.

<i>Table 39. Operating extremes with the power on</i>	
Measurement	Value
Altitude	0 - 2133 m (0 - 7000 ft)

<i>Table 39. Operating extremes with the power on (continued)</i>	
Measurement	Value
Dry bulb temperature	16 - 32°C (60 - 90°F)
Relative humidity	20 - 80%
Wet bulb temperature (maximum)	23°C (73°F)

Table 40 on page 80 provides the optimum operating points for your storage system with its power on.

<i>Table 40. Optimum operating points with the power on</i>	
Measurement	Value
Temperature	22°C (72°F)
Relative humidity	45%

Table 41 on page 80 provides the operating ranges for a storage system with power on.

<i>Table 41. Optimum operating ranges with power on</i>	
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 42 on page 80 provides the temperatures and humidity requirements for your storage system when the power is off.

<i>Table 42. Temperatures and humidity with the power off</i>	
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 43 on page 80 provides the temperatures and humidity requirements for storing your storage system.

<i>Table 43. Temperatures and humidity while in storage</i>	
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 44 on page 81](#). Additional information includes random vibration PSD profile breakpoints and operational shock levels. See [Table 45 on page 81](#) and [Table 46 on page 81](#).

Table 44. 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 45. 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 46. 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 it is determined that the levels of particulates or gases in your environment damaged the storage system, the warranty is void. Implementation of correctional measures is a customer responsibility.

The following criteria must be met:

Gaseous contamination

Severity level G1 according to ANSI/ISA 71.04-1985¹, which states that the reactivity rate of copper coupons must be fewer than 300 Angstroms per month ($\text{\AA}/\text{month}$, $\approx 0.0039 \mu\text{g}/\text{cm}^2\text{-hour weight gain}$)². In addition, the reactivity rate of silver coupons must be less than $300\text{\AA}/\text{month}$ ($\approx 0.0035 \mu\text{g}/\text{cm}^2\text{-hour weight gain}$)³. The reactive monitoring of gaseous corrosivity is conducted approximately 2 inches (5 cm) in front of the rack on the air inlet side at one-quarter and three-quarter frame height off the floor, or where the air velocity is much higher.

Particulate contamination

Data centers must meet the cleanliness level of ISO 14644-1 class 8. For data centers without airside economizers, the ISO 14644-1 class 8 cleanliness can be met by selecting one of the following filtration methods:

- The room air can be continuously filtered with MERV 8 filters.
- Air entering a data center can be filtered with MERV 11, or preferably MERV 13 filters.

For data centers with airside economizers, the choice of filters to achieve ISO class 8 cleanliness depends on the specific conditions present at that data center. The deliquescent relative humidity of the particulate contamination must be more than 60% RH⁴. Data centers must be free of zinc whiskers⁵.

1. ANSI/ISA-71.04.1985. *Environmental conditions for process measurement and control systems: Airborne contaminants*. Instrument Society of America, Research Triangle Park, NC, 1985.
2. The derivation of the equivalence between the rate of copper corrosion product thickness growth in $\text{\AA}/\text{month}$ and the rate of weight gain assumes that Cu_2S and Cu_2O grow in equal proportions.
3. The derivation of the equivalence between the rate of silver corrosion product thickness growth in $\text{\AA}/\text{month}$ and the rate of weight gain assumes that Ag_2S is the only corrosion product.
4. The deliquescent relative humidity of particulate contamination is the relative humidity at which the dust absorbs enough water to become wet and promote corrosion, ion migration, or both.
5. Surface debris is randomly collected from 10 areas of the data center on a 1.5 cm diameter disk of sticky, electrically conductive tape on a metal stub. If examination of the sticky tape in a scanning electron microscope reveals no zinc whiskers, the data center is considered free of zinc whiskers.

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
- Service navigation

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

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.

Providing service navigation equipment

You are responsible for providing appropriate service equipment for your storage system.

The DS8882F model 983 requires two service representatives to install at your site because of the weight of the components. You must provide the installation safety equipment based on the position of the components in your frame.

Table 47. Safety equipment.	
Feature code	Description
1101	5 ft. ladder for servicing components at a height of 35U or more
1102	3 ft. platform ladder for servicing components at a height of 20U or more
1103	Rolling step stool

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, transmit pertinent debugging data to IBM, and use remote code load as the delivery method for code upgrades.

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 communication.

The DS8882F uses remote code load (RCL) as the default delivery method for code upgrades. RCL provides the ability to securely download microcode and perform concurrent code load with support already engaged to quickly resolve any issues. RCL requires configuring Assist On-Site or remote support center (RSC), the same secure connection method used for the DS8000 system autonomic call home process for service. For on-site code load for license machine code upgrades, a separate contract is

required. Global Technology Services offers Global Total Microcode Support Services. For more information, contact your sales representative.

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 (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 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.

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 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 .

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 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 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 48 on page 92 compares data migration options and lists advantages and disadvantages of each.

Table 48. 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
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

Table 48. Comparison of data migration options (continued)

Type	Example	Advantages	Disadvantages
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.

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.

Note: DS8000 does not support communication with SKLM over IPP using custom device groups. KMIP is recommended for DS8000 systems communicating with IBM SKLM Key Servers in a Multi-Master configuration. When using IPP to communicate with IBM SKLM Key Servers in a Multi-Master

configuration, it is not possible for the DS8000 systems to automatically detect problems related to key redundancy, leaving you responsible for determining when high availability disaster recovery (HADR) synchronization is not functioning properly. Loss of data in the SKLM key store can result in loss of DS8000 data.

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 \(www.ibm.com/shop/publications/order\)](http://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 \(www.ibm.com/support/knowledgecenter/SSWPVP/\)](http://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 \(www-03.ibm.com/systems/services/labservices/platforms/labservices_storage.html\)](http://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 (`customer`).

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) (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 \(www.ibm.com/storage/dsfa\)](http://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 \(www.ibm.com/storage/dsfa\)](http://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 \(www.ibm.com/storage/dsfa\)](http://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](http://www.ibm.com/storage/dsfa) (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 \(www.ibm.com/storage/dsfa\)](http://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 \(www.ibm.com/storage/dsfa\)](http://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 \(www.ibm.com/storage/dsfa\)](http://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 \(www.ibm.com/support/knowledgecenter\)](http://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 \(www.ibm.com/able/\)](http://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 the eight modules that make up the DS8882F system. These modules are installed in a conforming 19-inch rack, in a 16U contiguous space, in the specified sequence (rack layout).

When the modules arrive, they contain any ordered I/O enclosures, device adapters, storage enclosures, drives, and the appropriate cables to support those components.

Power cords

Your shipment includes the power cords that you ordered.

Various media

The product 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. 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 87.

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

Appendix E. Compliance standards

This product was designed, tested, manufactured, and certified for safe operation. It complies with IEC 60950-1 and/or IEC 62368-1 and where required, to relevant national differences/deviations (NDs) to these IEC base standards. This includes, but is not limited to: EN (European Norms including all Amendments under the Low Voltage Directive), UL/CSA (North America bi-national harmonized and marked per accredited NRTL agency listings), and other such derivative certifications according to corporate determinations and latest regional publication compliance standardized requirements.

Regulatory Model ID (RMID, for control system base platform): AP2 / DS8882F

AP2 indicates the 2U Management Enclosure chassis only for regulatory component identity within the higher level integrated solution system being offered under the identity of DS8882F. As no rack or common cabinet enclosure is offered with this storage system of components, the MTM underneath this is a variable Print On Demand (Pod) based upon any given ordered integrated solution system, indicating it's service identity).

Note: Relevant Machine Type - Models (MT-Ms) may also be used to supplement identification (ID) for worldwide (WW) co-compliance filings or registrations with regulatory bodies.

Appendix F. European Union General Data Protection Regulation Considerations

General Data Protection Regulation (GDPR) compliance is required by the European Union (EU), effective May 25, 2018.

GDPR establishes a stronger data protection regulatory framework for processing customer data. GDPR impacts IBM and IBM's client contracts, policies, and procedures when handling personal data. GDPR features include:

- New and enhanced rights for individuals
- Widened definition of personal data
- New obligations for processors
- Potential for significant financial penalties for non-compliance
- Compulsory data breach notification

Read more about GDPR:

- <https://eugdpr.org>
- <https://www.ibm.com/data-responsibility/gdpr/>

Basic GDPR considerations

Complying with GDPR can present a challenge due to the variety of configurations and features available, and the number of ways products can be used separately and with third-party applications and systems. However, provides functions to help comply with GDPR, such as encryption of data at rest.

Your knowledge of Field Replaceable Units (FRUs) containing customer data is critical. The does not store or retain customer data in non-volatile memory other than data drives within the storage subsystem as written by your client applications. With encryption, if a drive fails and is replaced, the customer data is not readable from the replaced drive that is returned. Without encryption, you may need to retain and destroy failed drives to comply with GDPR. Your retention of the failed drive results in additional service costs due to the loss of failed parts credits.

Note: You are responsible for ensuring your compliance with various laws and regulations, including the EU GDPR. You are solely responsible for obtaining advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulations that may affect your business and any actions you may need to take to comply with such laws and regulations.

The products, services, and other capabilities described herein are not suitable for all client situations, and may have restricted availability. IBM does not provide legal, accounting, or auditing advice or represent or warrant that its services or products ensure that clients are in compliance with any law or regulation.

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 \(www.ibm.com/legal/copytrade.shtml\)](http://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

The *IBM Systems Safety Notices* 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 \(www.ibm.com/shop/publications/order/\)](http://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 was designed, tested, manufactured, and certified for safe operation. It complies with IEC 60950-1 and/or IEC 62368-1 and where required, to relevant National Differences/deviations (NDs) to these IEC base standards. This includes, but is not limited to: EN (European Norms including all Amendments under the Low Voltage Directive), UL/CSA (North America bi-national harmonized and marked per accredited NRTL agency listings), and other such derivative certifications according to corporate determinations and latest regional publication compliance standardized requirements.

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 123 and the [“Caution notices”](#) on page 117 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: If your system has a module containing a lithium battery, replace it only with the same module type made by the same manufacturer. The battery contains lithium and can explode if not properly used, handled, or disposed of.

Do not:

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

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

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

(R001 part 1 of 2)

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)

(R001 part 2 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 utstyr må kobles til jordside stikkontakter. PERIGO ALTA CORRENTE DE DISPERSÃO É essencial fazer ligação terra antes de conectar a ligação. FARA LACKSTRÖM Enheten måste vara jordad innan den ansluts. NEBEZPEČENSTVO MOŽNOST PRIENIKU NAPÁTIJA Uzemiť pred pripojením zdroja.	OPASNOST JAKO PROBLÉME STRUJE Uzemljenje je nužno prije spajanja dovoda napona. NEBEZPEČÍ VYSOKÝ URSÁKČÍ PROUD Pred pripojením napätia je nevyhnuté pripojiť zemný vodič. Fare! Høj ledstrøm Tilslutning til jord er nødvendig, før strømmen tilsluttes. GEVAAR HOOG LEKSTROOM Eerst aardverbinding 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 di protezione prima di collegare l'alimentazione. PERIGOR CORRENTE DE ALTA DISPERSIÓN Es imprescindible la conexión a tierra antes de conectar el suministro. VAARA KORkea VUOTOVIRTA Suojamaadoitus on välttämätön, ennen kuin laite liitetään sähköverkkoon. خطر فقدان التيار العالي الوصلة الأرضية ضرورية قبل توصيل التيار. KUVAUOZ YUVAHO PEYMA ALAPPOHE Ανατίττω: γείωση πριν την σύνδεση.	DANGER COURANT DE FUITE IMPORTANT Mise à la terre obligatoire avant le raccordement au réseau d'alimentation principale. Figyelem! NAGY SZIVARGÓ ÁRAM Földelés feltétlenül szükséges az üzembehelyezés előtt! PERIGOS ELEVADA CORRENTE DE FUGA Imprescindível a ligação a terra antes da conexão do alimentador. OPACHO BYSOKÝ TOK VTEVOKI Перед подключением питания необходимо заземление. NEBEZPEČENSTVO ZNACZNY PRĄD UPŁYWU Przed podłączeniem zasilania należy konieczne podłączyć uziemienie.	危險 嚴重漏電 連接電源前，必須接地。 高漏電流 連接前，務必先完成接地。 危險 大漏電流 電源を接続する前に必ず接地してください。 위험 노출 위험 전원 공급전배 전도체 접지 연결을. סכנה ודם ולינה נכונה חיבור לזרמת חדרתי לפני חיבור למקור הכוח.
---	---	---	---	---

and / or



Vendor-specific uninterruptible power supply safety statements

These statements apply specifically to the vendor uninterruptible power supply (UPS).

General caution statements

- Do not use extension cords to connect to an AC outlet.
- Do not plug the UPS's input cord into its own output receptacles.
- Before using the UPS, you must unpack the UPS and allow it to adjust to room temperature (20-25degC) for at least two hours to avoid moisture condensing inside the UPS.
- Do not connect any equipment that may overload the UPS.
- Contact qualified service personnel if either of the following events occur:
 - Liquid is poured or splashed on the UPS.
 - The UPS does not run normally after its user manual has been carefully observed.
- The allowable storage temperature and relative humidity (non-condensing) are -15..+50 C and 5..95%RH.

Vendor-specific warning, caution, and danger statements

Warning: To reduce the risk of fire or electric shock, install in a temperature and humidity-controlled indoor area, free of conductive contaminants. Refer servicing to qualified service personnel only. Intended for installation in a controlled environment.



CAUTION:

Risk of electrical shock. Capacitor stores hazardous energy. Do not remove cover for 5 minutes after disconnecting all sources of supply.

Risk of electric shock. Hazardous live parts inside this unit are energized from the battery supply even when the input AC power is disconnected.

Risk of electric shock. Do not remove cover. No user serviceable parts inside. Refer servicing to qualified service personnel.

Risk of electric shock. This UPS receives power from more than one source. Disconnection of the AC sources and the DC sources is required to de-energize this unit before servicing.

Warning: Turn OFF the UPS and also cut off the AC input source before performing battery pack replacement.



CAUTION:

Risk of electrical shock. This battery cabinet receives power from more than one source.

The lead acid battery may cause chemical hazard.

This battery presents a risk of electric shock and energy hazard.

Lead (Pb) acid battery in the inside of the enclosure.

For disposal instructions for the battery, see user's manual.

Only open the battery compartment in the described way, do not disassemble other parts of the UPS.



DANGER: Risk of electric shock. Do not touch uninsulated battery terminal.



CAUTION:

Safety instructions - save these instructions! This manual contains important instructions that should be followed during installation and maintenance of the UPS and batteries.

This UPS contains batteries, which are potentially hazardous to user, even when the UPS is not connected to the utility power.

The AC output of the UPS need a disconnect switch such as a breaker which has to be provided by users.

The over-current protection for the output AC circuit has to be provided by users. The UPS has an electronic protection of AC output short circuit.

Servicing of batteries should be performed or supervised by personnel who are knowledgeable about batteries and the required precautions. Keep unauthorized personnel away from batteries.

When replacing batteries, replace with the same number and type of battery.

CAUTION! Do not dispose of battery or batteries in a fire. The batteries may explode.

CAUTION! Do not open or mutilate the battery or batteries. The released electrolyte is harmful to the skin and eyes. It may be toxic.

CAUTION! A battery can present a risk of electrical shock and high short circuit current. The following precautions should be observed when working on batteries.

- Remove watches, rings, or other metal objects.
- Use tools with insulated handles.
- Wear rubber gloves and boots.
- Do not lay tools or metal parts on the top of batteries.
- Disconnect charging source prior to connecting or disconnecting battery terminals.
- Determine if the battery is inadvertently grounded. If inadvertently grounded, remove source of ground. Contact with any part of a grounded battery can result in electrical shock. The likelihood of such shock will be reduced if such grounds are moved during installation and maintenance (applicable to UPS and a remote battery supply not having a grounded supply circuit).

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

RUSSIAN

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。

taitemi

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

12c00790

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 [24](#)
3.5-inch [24](#)
4-port HA [25](#)
8-port HA [25](#)

A

accessibility
 features [103](#)
acclimation [73](#)
activating licenses [100](#)
adapters
 configuration rules [56](#)
 Fibre Channel host [55](#)
air circulation [79](#)
air quality [81](#)
algorithms [29](#)
allocation methods [14](#)
AS/400 LUN
 control switch settings [89](#)
Attachment to IBM Z [89](#)
auto-rebalance [34](#)
auxiliary volumes [34](#)
availability features [24](#)

B

balancing the I/O load [26](#)
Base function license
 Database Protection [66](#)
 Easy Tier [67](#)
 Encryption Authorization [67](#)
 I/O Priority Manager [67](#)
 Operating environment license [67](#)
Base Function license
 thin provisioning [67](#)
battery backup modules
 feature codes [61](#)
BSMI certificate [62](#)
BTU [78](#)

C

cable
 feature codes, Fibre Channel cable [55](#)
 Fibre Channel host adapter [55](#)
 I/O adapters [54](#)
cables
 drive [52](#)
cache [57](#)
capacity
 calculating physical and effective [52](#)
 exhaust [78](#)
caution notices [116](#), [117](#)

CCW, channel command words [27](#)
certificate, BSMI [62](#)
China non-encryption certificate [62](#)
circuit breakers
 high-voltage [78](#)
 low-voltage [78](#)
CKD, count key data storage [27](#)
CLI, command-line interface [16](#)
cloud
 transparent cloud tiering [31](#)
clusters, RAID disk groups [26](#)
CoD [24](#)
cold weather procedures [73](#)
comments, sending [xii](#)
communication requirements, host attachment [84](#)
company information [87](#)
configuration
 battery backup modules [61](#)
 processor memory [57](#)
configuration control indicators [49](#)
configuration overview [9](#)
Configuration overview [11](#)
configuration rules
 host adapters [56](#)
 Standby CoD disk drives [52](#)
 storage devices [52](#)
 storage enclosures [52](#)
connectivity
 I/O enclosures [54](#)
consolidating storage [27](#)
containers, shipping [73](#)
contamination information [81](#)
control switch settings [89](#)
control unit threshold [89](#)
conventions
 terminology [vii](#)
 typefaces [vii](#)
cooling [79](#)
Copy Services
 considerations [37](#)
 disaster recovery [44](#)
 licensed functions [46](#)
 overview [37](#)
 Safeguarded copy function [70](#)
Copy Services license
 FlashCopy function [70](#)
 Remote mirror and copy [70](#)
 z/OS Global Mirror [71](#)
 z/OS Metro/Global Mirror Incremental Resync [71](#)
Copy Services Manager [17](#)
Copy Services Manager on hardware management console
 license [71](#)
Copy Services software package
 Remote mirror and copy [70](#)
corrosive gasses and particulates [81](#)
count key data storage [27](#)
CUIR, control-unit initiated reconfiguration [89](#)

customization worksheets [109](#)

D

danger notices [116](#), [123](#)

data

securing [48](#)

data migration

selecting method [92](#)

data placement [26](#)

Database Protection [66](#)

description

VMware [35](#)

description of Easy Tier [34](#)

description of EAV [30](#)

Device threshold [89](#)

dimensions

storage system, installed [76](#)

disaster recovery

Copy Services [44](#)

disk drive

cable [50](#)

disk drive module

maintenance policy [25](#)

disk drive sets [50](#)

disk drives [24](#)

disk enclosures

fillers [50](#)

disk intermix

configuration indicators [49](#)

Disk Manager

monitoring performance [17](#)

Disk Storage Feature Activation (DSFA) [100](#)

drive set

capacity [52](#)

drive sets [51](#)

drives

cable [52](#)

capacity calculation [52](#)

DS command-line interface [16](#)

DS8000 Storage Management GUI [15](#)

DSFA, Disk Storage Feature Activation [100](#)

dynamic expansion

volume [32](#)

dynamic volume expansion [32](#)

E

EAM [28](#)

Easy Tier

overview [34](#)

EAV CKD

1 TB IBM Z CKD [30](#)

3390 Model A [30](#)

cylinder size [30](#)

enclosure fillers

feature codes [51](#)

encryption

overview [48](#)

planning for [95](#), [96](#)

Encryption Authorization [67](#)

encryption, disabling [62](#)

environment

environment (*continued*)

operating requirements [79](#)

environmental notices [127](#)

ESE capacity controls [33](#)

ESE volumes [34](#)

exhaust [78](#)

expansion model position

configuration indicators [49](#)

extended address volumes

overview [30](#)

Extended Remote Copy (XRC) (see z/OS Global Mirror) [71](#)

F

failover and fallback [44](#)

FB, fixed block [27](#)

FDE [24](#)

Feature Activation (DSFA), IBM Disk Storage [100](#)

feature codes

additional setup options [61](#)

battery backup modules [61](#)

drive sets [51](#)

enclosure fillers [51](#)

features, other configuration [61](#)

Fibre Channel cable [55](#)

Fibre Channel host adapters [55](#)

flash drives [51](#)

I/O adapter [54](#)

I/O enclosure [54](#)

management console [50](#)

memory [57](#)

optional [61](#)

ordering optional features [49](#)

PCIe adapters [55](#)

physical configuration [49](#)

power cords [58](#)

power features [61](#)

processors [57](#)

setup [61](#)

shipping [61](#)

shipping and setup [61](#)

Transparent cloud tiering adapters [57](#)

features

input voltage

about [61](#)

feedback, sending [xii](#)

Fibre Channel

host adapters [55](#)

host adapters feature codes [55](#)

host attachment [84](#)

open-systems hosts [25](#)

SAN connections [25](#)

Fibre Channel cable

feature codes [55](#)

fire suppression [82](#)

fixed block storage [27](#)

flash copy [45](#)

flash drives [24](#), [51](#)

flash drives, feature codes [51](#)

FlashCopy

Multiple incremental FlashCopy [39](#)

floor and space requirements [76](#)

force option [28](#), [33](#)

G

GDPR [113](#)
General Data Protection Regulation [113](#)
Global mirror [70](#)
Global Mirror [44](#), [45](#)

H

HA intermix [25](#)
hardware features [21](#)
hardware planning [21](#)
High Performance FICON [68](#)
high-voltage
 installations [78](#)
homogeneous pools [34](#)
homologation [116](#)
host adapters
 Fibre Channel [55](#)
host attachment
 overview [25](#)
host systems
 communication requirements [84](#)
hot spot management [34](#)
how to order using feature codes [49](#)
HyperPAV [68](#)

I

I/O adapter
 features [54](#)
I/O enclosure
 feature codes [54](#)
I/O enclosures [54](#)
I/O load, balancing [26](#)
I/O plug order [25](#)
I/O Priority Manager [67](#)
IBM Copy Services Manager
 copy services [17](#)
 Replication [17](#)
IBM Disk Storage Feature Activation (DSFA) [100](#)
IBM Spectrum Control [17](#)
IBM Z
 HyperPAV [68](#)
 parallel access volume [68](#)
 Parallel Access Volumes [68](#)
 power control settings [89](#)
IBM Z hosts
 FICON attachment overview [25](#)
IEC 60950-1 [116](#)
implementation, RAID [13](#)
initialization [30](#)
input voltage
 configuration rules [61](#)
input voltage requirements [78](#)
installation
 components with shipment [107](#)
installation planning [73](#)
IOPS [34](#)

L

labels [125](#)

labels, safety information [116](#)
LFF [24](#)
licensed functions
 Copy Services [46](#)
licenses
 Disk Storage Feature Activation (DSFA) [100](#)
 function authorization documents [100](#)
logical subsystems
 overview [14](#)
logical volumes [28](#)
low-voltage
 installations [78](#)
LSS, logical subsystem [14](#)
LUN
 calculation [29](#)
 control switch settings [89](#)

M

machine types [2](#)
maintenance policy
 disk drive module [25](#)
managed allocation [14](#)
management console
 ESSNI [50](#)
 feature codes [50](#)
 HMC [50](#)
 HMC remote access [50](#)
 multiple storage systems [24](#)
 network settings [87](#)
 overview [24](#)
 TCP/IP ports [83](#)
management consoles
 feature codes [50](#)
management interfaces [15](#)
Media threshold [89](#)
memory
 feature codes [57](#)
 system memory [57](#)
Metro mirror [70](#)
Metro/Global mirror [70](#)
migrating data
 selecting method [92](#)
Multiple incremental FlashCopy [39](#)

N

network settings [87](#)
new features [xiii](#)
nodes [84](#)
noise level [78](#)
non-encryption certificate [62](#)
notices
 caution [117](#)
 danger [123](#)
 environmental [127](#)
 labels [125](#)
 safety [116](#)
notification settings
 methods [88](#)

O

- obtaining activation codes [100](#)
- operating environment
 - power on [79](#)
 - while in storage [79](#)
 - with power on or off [79](#)
- Operating environment license [67](#)
- overview
 - host attachment [25](#)

P

- Parallel Access Volumes
 - understanding static and dynamic [36](#)
- pass-through [45](#)
- PAV (Parallel Access Volumes) [36](#)
- PCIe
 - adapter feature codes [55](#)
- performance gathering
 - Disk Manager [17](#)
- physical configuration
 - drive capacity [52](#)
 - I/O adapter features [54](#)
 - I/O cable [54](#)
 - I/O enclosures [54](#)
 - input voltage of power supply [61](#)
 - management consoles [50](#)
 - power cords [58](#)
 - power features [61](#)
 - processors [57](#)
- planning
 - activating
 - full-disk encryption [96](#)
 - disk encryption
 - activating [96](#)
 - planning [96](#)
 - encryption [95, 96](#)
 - environment requirements [79](#)
 - full-disk encryption activation [96](#)
 - IBM Security Key Lifecycle Manager [95, 96](#)
 - IBM Storage Appliance 2421 Model AP1 [96](#)
 - isolated key servers [96](#)
 - network and communications [83](#)
 - operating environment, power [80](#)
 - power connector [78](#)
 - safety [82](#)
 - storage complex setup [87](#)
- point-in-time copy [70](#)
- pool rebalancing [34](#)
- power
 - consumption [78](#)
 - operating environment, off [80](#)
 - operating environment, on [79](#)
 - outlet requirements [77](#)
- power connector
 - requirements [78](#)
 - specifications [78](#)
- power control settings [89](#)
- power cords
 - feature codes [58](#)
 - power connector [78](#)
- power features
 - configuration rules [61](#)

- power frequencies [78](#)
- power supply
 - input voltage of [61](#)
- Present SIM data to all hosts [89](#)
- processor
 - feature codes [57](#)
 - memory (cache) [57](#)
- processors
 - feature codes [57](#)
- publications
 - ordering [xi](#)
 - product [vii](#)
 - related [vii](#)

Q

- quick initialization [30](#)

R

- rack requirements [74](#)
- RAID
 - disk groups [26](#)
 - implementation [13](#)
 - RAID 10 overview [13](#)
 - RAID 5 overview [13](#)
 - RAID 6 overview [13](#)
 - RAID overview [13](#)
- rank depopulation [34](#)
- redundancy
 - management consoles [50](#)
- Remote mirror and copy [70](#)
- remote mirror for z Systems (see z/OS Global Mirror) [71](#)
- remote power control [84](#)
- remote support
 - connections [83](#)
 - settings [88](#)
- replication
 - copy services functions [17](#)
- requirements
 - floor and space [76](#)
 - floor load [76](#)
 - host attachment communication [84](#)
 - input voltage [78](#)
 - loading dock [74](#)
 - planning network and communications [83](#)
 - power connectors [78](#)
 - power outlets [77](#)
 - rack [74](#)
 - receiving area [74](#)
 - space [76](#)
- resource groups
 - copy services [45](#)
- RGL [45](#)
- rotate capacity [14](#)
- rotate volumes [14](#)
- Russia non-encryption certificate [62](#)

S

- safety
 - fire suppression [82](#)
 - information labels [116](#)

- safety (*continued*)
 - notices [116](#)
 - operating environment [82](#)
 - power outlets [77](#)
 - service navigation [83](#)
 - temperature and cooling [82](#)
- SAN
 - connections with Fibre Channel adapters [25](#)
- SAS [24](#)
- SAS enterprise and NL SAS [34](#)
- SATA [24](#)
- scenarios
 - adding storage [101](#)
- scope limiting
 - disaster recovery [45](#)
- security
 - best practices
 - service accounts [97](#)
 - user accounts [97](#)
- serial number setting [89](#)
- Service Information Messages [89](#)
- service navigation [83](#)
- SFF [24](#)
- shipments
 - authorized service components [108](#)
 - container weight, dimensions [73](#)
 - hardware, software [107](#)
 - loading ramp [74](#)
 - media [108](#)
 - planning for receipt [74](#)
 - receiving area [74](#)
 - requirements
 - loading ramp [74](#)
- shipping containers [73](#)
- SIM [89](#)
- slot plug order [25](#)
- space requirements [76](#)
- specifications
 - power connectors [78](#)
- standards
 - air quality [81](#)
- Standby CoD disk drives [52](#)
- statement of limited warranty [105](#)
- storage area network
 - connections with Fibre Channel adapters [25](#)
- storage features
 - configuration rules [52](#)
- storage system
 - architecture [3](#)
 - implementation [3](#)
- storage-enclosure fillers [51](#)
- storage, consolidating [27](#)
- System i
 - control switch settings [89](#)

T

- T10 DIF
 - ANSI support [27](#)
 - Data Integrity Field [27](#)
 - FB LUN [27](#)
 - FB LUN protection [27](#)
 - FB volume [27](#)
 - Linux on z Systems [27](#)

- T10 DIF (*continued*)
 - SCSI end-to-end [27](#)
 - standard protection [27](#)
 - z Systems [27](#)
- Taiwan BSMI certificate [62](#)
- terminology [vii](#)
- thermal load [78](#)
- thin provisioning
 - ESE capacity controls [33](#)
 - overview [33](#)
- thin provisioning and Easy Tier [34](#)
- three tiers [34](#)
- Tivoli Storage Productivity Center [24](#)
- trademarks [116](#)
- transparent cloud tiering [31, 69](#)
- Transparent cloud tiering
 - adapters feature codes [57](#)

U

- understanding fixed block (FB) architecture [27](#)
- understanding logical volumes [28](#)
- user interfaces [15](#)

V

- VMware
 - Array Integration support [35](#)
 - restrictions [35](#)
- volume capacity
 - overview [29](#)
- volume deletion
 - force option [33](#)
 - safe option [33](#)
- volume migration [34](#)
- volume rebalancing [34](#)
- volumes
 - allocation [14, 28](#)
 - deletion [28](#)
 - force option [28](#)
 - modification [28](#)

W

- warranty [105](#)
- websites [vii](#)
- weight
 - storage system, installed [76](#)
- weight and dimensions
 - shipping container [73](#)
- worksheets
 - provided equipment [107](#)
 - customization [109](#)
- WWID, switch settings [89](#)

X

- XRC (Extended Remote Copy) (see z/OS Global Mirror) [71](#)

Z

- z-synergy Services license
 - High Performance FICON [68](#)

z-synergy Services license (*continued*)

HyperPAV [68](#)

Parallel Access Volumes [68](#)

transparent cloud tiering [69](#)

z/OS Distributed Data Backup [69](#)

z/OS Global Mirror [71](#)

z/OS Metro/Global Mirror Incremental Resync [71](#)



Printed in USA

GC27-9259-02

