

IBM Storage DS8900F Architecture and Implementation Updated for Release 9.3.2

Peter Kimmel

Daniel Beukers

Jeff Cook

Bozhidar Feraliev

Jörg Klemm

Connie Riggins

Gaurav Sabharwal



Storage



IBM Redbooks

**IBM Storage DS8900F Architecture and
Implementation: Updated for Release 9.3.2**

April 2023

Note: Before using this information and the product it supports, read the information in “Notices” on page xi.

Fourth Edition (April 2023)

This edition applies to DS8900F systems with IBM Storage DS8000 Licensed Machine Code (LMC) 7.9.32 (bundle version 89.32.xx.x), referred to as Release 9.3.2.

© Copyright International Business Machines Corporation 2020, 2023. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	xi
Trademarks	xii
Preface	xiii
Authors	xiv
Now you can become a published author, too!	xv
Comments welcome	xv
Stay connected to IBM Redbooks	xvii
Part 1. Concepts and architecture	1
Chapter 1. Introducing the IBM DS8900F storage system	3
1.1 Introduction to the DS8900F system	4
1.1.1 At-a-glance features of the DS8900F	7
1.2 DS8900F controller options and frames	11
1.3 DS8900F architecture and functions overview	12
1.3.1 Overall architecture and components	12
1.3.2 Storage capacity	15
1.3.3 Supported environments	15
1.3.4 Configuration flexibility	16
1.3.5 Copy Services functions	18
1.3.6 Service and setup	19
1.3.7 IBM Certified Secure Data Overwrite	20
1.4 Performance features	21
1.4.1 4.3 TB cache	21
1.4.2 32 Gbps and 16 Gbps host adapters	21
1.4.3 Sophisticated caching algorithms	21
1.4.4 Flash storage	22
1.4.5 Performance for IBM Z	23
1.4.6 Performance enhancements for IBM Power servers	24
Chapter 2. IBM DS8900F hardware components and architecture	25
2.1 Flash drive terminology of the DS8900F	26
2.1.1 Storage system	26
2.1.2 Management Enclosure	27
2.1.3 Central processor complex	27
2.2 DS8900F configurations and models	28
2.2.1 DS8980F Analytic Class configuration	28
2.2.2 DS8950F Agility Class configuration	30
2.2.3 DS8910F Flexibility Class Racked configuration	31
2.2.4 DS8910F Flexibility Class Rack-Mounted configuration	32
2.2.5 DS8900F base frames	35
2.2.6 DS8900F expansion frame	38
2.2.7 Scalable upgrades	41
2.2.8 Licensed functions	41
2.3 DS8900F architecture overview	42
2.3.1 IBM POWER9 processor-based CPCs	42
2.3.2 Processor memory	45
2.3.3 Flexible service processor	46

2.3.4	Ethernet connections	46
2.3.5	Peripheral Component Interconnect Express adapters	47
2.4	I/O enclosures and adapters	48
2.4.1	Cross-cluster communication	51
2.4.2	I/O enclosure adapters	51
2.5	Flash drive enclosures	56
2.6	Power and cooling	58
2.6.1	Rack power control cards	58
2.6.2	Intelligent Power Distribution Units	59
2.6.3	Power domains	64
2.6.4	Rack management 24-port Ethernet switch pair	64
2.6.5	Backup Power Modules and NVDIMM	65
2.6.6	Power cord options	66
2.6.7	Enclosure power supply units	66
2.7	Management Console and network	67
2.7.1	Ethernet switches	68
Chapter 3. IBM DS8900F reliability, availability, and serviceability		71
3.1	DS8900F processor complex features	72
3.1.1	POWER9 PowerVM Hypervisor	72
3.1.2	POWER9 processor	72
3.1.3	Cross-cluster communication	76
3.1.4	Environmental monitoring	76
3.1.5	Resource deconfiguration	77
3.2	CPC failover and failback	78
3.2.1	Dual cluster operation and data protection	78
3.2.2	Failover	79
3.2.3	Failback	80
3.2.4	NVS and power outages	81
3.3	Data flow in the DS8900F	82
3.3.1	I/O enclosures	82
3.3.2	Host connections	82
3.3.3	Metadata checks	86
3.4	RAS on the Hardware Management Console	86
3.4.1	Licensed Internal Code updates	87
3.4.2	Call Home and remote support	87
3.5	RAS on the storage system	88
3.5.1	RAID configurations	89
3.5.2	Drive path redundancy	90
3.5.3	Flash RAID controller redundancy	90
3.5.4	Predictive failure analysis	91
3.5.5	Disk scrubbing	91
3.5.6	RAID support	91
3.5.7	RAID 6 overview	92
3.5.8	RAID 10 overview	94
3.5.9	RAID 5 implementation in DS8900F	95
3.5.10	Smart Rebuild	95
3.5.11	Spare creation	96
3.6	RAS on the power subsystem	97
3.6.1	Power components	97
3.6.2	Line power loss	99
3.6.3	Line power fluctuation	99
3.6.4	Power control	99

3.7 Other features	101
3.7.1 Internal network	101
3.7.2 Earthquake resistance	101
3.7.3 IBM Certified Secure Data Overwrite	101
Chapter 4. Virtualization concepts	107
4.1 Virtualization definition	108
4.2 Benefits of virtualization	108
4.3 Abstraction layers for drive virtualization.	109
4.3.1 Array sites	109
4.3.2 Arrays	109
4.3.3 Ranks	111
4.4 Extent pools	113
4.4.1 Dynamic extent pool merge	114
4.4.2 Logical volumes	116
4.4.3 Allocating, deleting, and modifying LUNs and CKD volumes	121
4.4.4 Volume allocation and metadata.	125
4.4.5 Logical subsystems.	130
4.4.6 Volume access	133
4.4.7 Virtualization hierarchy summary	135
4.5 Terminology for IBM Storage products	137
Part 2. Planning and installation	139
Chapter 5. IBM DS8900F physical planning and installation	141
5.1 Considerations before the installation: Planning for growth	142
5.1.1 Client responsibilities for the installation.	142
5.1.2 Participants	143
5.1.3 Required information.	143
5.2 Planning for the physical installation	144
5.2.1 Delivery and staging area	144
5.2.2 Floor type and loading	145
5.2.3 Overhead cabling features	146
5.2.4 Room space and service clearance	147
5.2.5 Power requirements and operating environment	148
5.2.6 Host interface and cables	150
5.2.7 Host adapter Fibre Channel specifics for open environments	153
5.2.8 FICON specifics on a z/OS environment	154
5.2.9 Best practices for host adapters	154
5.2.10 Worldwide node name and worldwide port name determination	155
5.3 Network connectivity planning.	159
5.3.1 Hardware Management Console and network access	159
5.3.2 IBM Spectrum Control and IBM Storage Insights	160
5.3.3 DS Command-Line Interface.	160
5.3.4 Remote support connection	161
5.3.5 Storage area network connection	161
5.3.6 Key manager servers for encryption.	161
5.3.7 Lightweight Directory Access Protocol server.	163
5.4 Remote Mirror and Remote Copy connectivity	163
5.5 Disk capacity considerations.	163
5.5.1 Disk sparing	164
5.5.2 Disk capacity	164
Chapter 6. IBM DS8900F Management Console planning and setup.	167

6.1 DS8900F Management Console overview	168
6.1.1 Management Enclosure	168
6.1.2 Management Console hardware	170
6.1.3 Private and Management Ethernet networks	171
6.2 Management Console software	174
6.2.1 DS Storage Management GUI	175
6.2.2 Data Storage Command-Line Interface	175
6.2.3 RESTful application programming interface	175
6.2.4 IBM Copy Services Manager interface	175
6.2.5 Updating the embedded IBM Copy Services Manager	176
6.2.6 Web User Interface	180
6.2.7 IBM ESSNI server	182
6.3 Management Console activities	182
6.3.1 Management Console planning tasks	182
6.3.2 Planning for Licensed Internal Code upgrades	183
6.3.3 Time synchronization	184
6.3.4 Monitoring DS8900F with the Management Console	184
6.3.5 Event notification through syslog	184
6.3.6 Call Home and remote support	185
6.4 Management Console network settings	185
6.4.1 Private networks	186
6.5 User management	187
6.5.1 Password policies	187
6.5.2 Remote authentication	188
6.5.3 Service Management Console User Management	189
6.5.4 Service Management Console LDAP authentication	195
6.5.5 Multifactor authentication (MFA)	196
6.6 Secondary Management Console	197
6.6.1 Management Console redundancy benefits	197
Chapter 7. IBM DS8900F features and licensed functions	199
7.1 DS8900F licensed functions	200
7.1.1 General introduction to licensing	202
7.1.2 Licensing cost structure	204
7.2 Activating licensed functions	207
7.2.1 Obtaining DS8000 machine information and activating license keys	207
7.2.2 Obtaining the activation codes	213
7.2.3 Applying activation codes by using the DS CLI	218
7.3 Licensed scope considerations	220
7.4 Expert Care	220
Part 3. Storage configuration	223
Chapter 8. Configuration flow	225
8.1 Configuration worksheets	226
8.2 User and role management	226
8.3 Encryption	228
8.4 Network security	228
8.5 Configuration flow	229
8.6 General storage configuration guidelines	230
Chapter 9. IBM DS8900F Storage Management GUI	233
9.1 Introduction	234
9.2 DS GUI: Getting started	235

9.2.1	Accessing the DS GUI	235
9.2.2	System Setup wizard	236
9.2.3	Configuring Fibre Channel port protocols and topologies	239
9.3	Managing and monitoring the storage system	241
9.3.1	Storage Monitoring and Servicing from the Unified Service GUI	248
9.3.2	Storage Management help functions	254
9.4	System configuration overview	254
9.4.1	Network settings	254
9.4.2	Security settings	256
9.4.3	System Settings	260
9.4.4	Notifications settings	264
9.4.5	Support settings	267
9.5	Logical configuration overview	270
9.6	Logical configuration for open systems volumes	271
9.6.1	Configuration flow	271
9.6.2	Creating FB pools for open systems hosts	271
9.6.3	Creating FB volumes for open systems hosts	278
9.6.4	Creating FB host attachments	282
9.6.5	Assigning FB volumes	291
9.7	Logical configuration for Count Key Data volumes	292
9.7.1	Configuration flow	292
9.7.2	Creating CKD storage pools	293
9.7.3	Creating CKD logical subsystems	296
9.7.4	Creating CKD volumes	299
9.7.5	Creating CKD parallel access volumes	300
9.7.6	Setting the FC port protocols for IBM Z attachment	304
9.8	Expanding volumes	304
9.9	Deleting a pool	306
9.10	Deleting volumes	306
9.11	Reinitializing a thin-provisioned volume	308
9.12	Easy Tier support	308
9.13	Monitoring system health	309
9.13.1	Hardware components: Status and attributes	311
9.13.2	Viewing components health and state from the system views	317
9.13.3	Monitoring system events	319
9.13.4	Exporting system-wide information	319
9.13.5	Audit logs	321
9.14	Performance monitoring	322
9.14.1	Performance statistics	323
9.14.2	Working with customized performance graphs	326
9.15	Fibre Channel error rate statistics	336
9.16	Providing feedback	337
Chapter 10.	IBM DS8900F Storage Management Command-line Interface	339
10.1	DS CLI overview	340
10.1.1	Supported operating systems for the DS CLI	340
10.1.2	Installation hints and tips	341
10.1.3	Installing the DS CLI on a Windows 10 or 11 system	342
10.1.4	Installing the DS CLI on an z/OS system	342
10.1.5	DS CLI version	347
10.1.6	User accounts	348
10.1.7	User management by using the DS CLI	348
10.1.8	DS CLI profile	350

10.1.9	Configuring the DS CLI to use the second HMC	353
10.1.10	Command structure	353
10.1.11	Using the DS CLI application	353
10.1.12	Return codes	356
10.1.13	User assistance	356
10.2	I/O port configuration	357
10.3	DS8900F storage configuration for Fixed-Block volumes	358
10.3.1	Disk classes	358
10.3.2	Creating the arrays	358
10.3.3	Creating the ranks	361
10.3.4	Creating the extent pools	362
10.3.5	Creating the FB volumes	363
10.3.6	Creating the volume groups	370
10.3.7	Creating host connections and clusters	372
10.3.8	Mapping open system host disks to storage unit volumes	380
10.4	DS8900F storage configuration for the CKD volumes	380
10.4.1	Creating the arrays	381
10.4.2	Creating the ranks and extent pools	381
10.4.3	Logical control unit creation	382
10.4.4	Creating the CKD volumes	383
10.4.5	Resource groups	390
10.4.6	IBM Easy Tier	390
10.5	Metrics with DS CLI	391
10.5.1	Offload performance data and other parameters	397
10.6	Private network security commands	398
10.7	Copy Services commands	400
10.8	Earlier DS CLI commands and scripts	401
10.9	For more information	402

Part 4. Maintenance and upgrades 403

Chapter 11. Licensed Machine Code	405	
11.1	How new Licensed Internal Code is released	406
11.2	Bundle installation	408
11.2.1	Remote Code Load	412
11.2.2	Customer Code Load	413
11.3	Code updates	418
11.4	Host adapter firmware updates	418
11.4.1	Light-on fastload firmware update	419
11.4.2	Remote Mirror and Remote Copy path considerations	419
11.4.3	Control-unit initiated reconfiguration	419
11.5	Loading the code bundle	419
11.6	Fast path concurrent code load	420
11.7	Postinstallation activities	421
11.8	Summary	421
Chapter 12. Monitoring and support	423	
12.1	SNMP implementation on the DS8900F	424
12.1.1	Message Information Base file	424
12.1.2	Predefined SNMP trap requests	424
12.2	SNMP notifications	425
12.2.1	Serviceable event that uses specific trap 3	425
12.2.2	Copy Services event traps	426
12.2.3	Thin-provisioning SNMP	431

12.3	SNMP configuration	432
12.3.1	SNMP preparation	432
12.3.2	SNMP configuration with the HMC	432
12.3.3	SNMP configuration with the DS CLI	436
12.4	Introducing remote support	436
12.5	IBM policies for remote support	437
12.6	Remote support advantages	437
12.7	Remote support and Call Home	437
12.7.1	Call Home and heartbeat: Outbound	437
12.7.2	Data offload: Outbound	438
12.7.3	Outbound connection types	438
12.8	Remote Support Access (inbound)	439
12.8.1	Assist On-site	439
12.8.2	DS8900F-embedded AOS	440
12.8.3	IBM Remote Support Center for DS8900F	441
12.8.4	Support access management through the DS CLI and DS GUI	441
12.9	Call Home and Assist On-Site customer-provided certificates	444
12.10	Audit logging	446
12.11	Using IBM Storage Insights	448
12.11.1	IBM Storage Insights	448
12.11.2	Getting started with IBM Storage Insights	449
12.11.3	Case interaction with IBM Storage Insights	452
12.11.4	IBM Storage Insights: Alert Policies	459
12.12	IBM Call Home Connect Cloud	460
	Abbreviations and acronyms	463
	Related publications	467
	IBM Redbooks	467
	Other publications	468
	Online resources	468
	Help from IBM	469

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

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 and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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 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.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, 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 web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®	IBM Research®	PowerPC®
C3®	IBM Security®	PowerVM®
Db2®	IBM Services®	RACF®
DS8000®	IBM Spectrum®	Redbooks®
Easy Tier®	IBM Z®	Redbooks (logo)  ®
Enterprise Storage Server®	IBM z13®	WebSphere®
eServer™	IBM z14®	z/Architecture®
FICON®	IBM z16™	z/OS®
FlashCopy®	Interconnect®	z/VM®
GDPS®	Parallel Sysplex®	z13®
Guardium®	PIN®	z15®
HyperSwap®	POWER®	z16™
IBM®	Power Architecture®	zEnterprise®
IBM Cloud®	POWER8®	
IBM FlashSystem®	POWER9™	

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, and the Windows logo 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.

OpenShift, Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

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

VMware, VMware vSphere, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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

Preface

This IBM Redbooks publication describes the concepts, architecture, and implementation of the IBM Storage DS8900F family. The book provides reference information to assist readers who need to plan for, install, and configure the DS8900F systems. This edition applies to DS8900F systems with IBM Storage DS8000® Licensed Machine Code (LMC) 7.9.32 (bundle version 89.32.xx.x), referred to as Release 9.3.2.

This book was updated in April 2023 to include the following enhancements that were delivered with DS8000 Release 9.3.2:

- ▶ Multifactor authentication (MFA)
- ▶ Call Home and Assist On-Site customer-provided certificates
- ▶ Support remote code load (RCL) for IBM Expert Care Advanced clients (Expert Care)

The DS8900F systems are all-flash exclusively, and they are offered as three classes:

- ▶ DS8980F: Analytic Class

The DS8980F Analytic Class offers best performance for organizations that want to expand their workload possibilities to artificial intelligence (AI), Business Intelligence (BI), and machine learning (ML).

- ▶ IBM DS8950F: Agility Class

The Agility Class consolidates all your mission-critical workloads for IBM Z®, IBM LinuxONE, IBM Power, and distributed environments under a single all-flash storage solution.

- ▶ IBM DS8910F: Flexibility Class

The Flexibility Class reduces complexity while addressing various workloads at the lowest DS8900F family entry cost.

The DS8900F architecture relies on powerful IBM POWER9™ processor-based servers that manage the cache to streamline disk input/output (I/O), which maximizes performance and throughput. These capabilities are further enhanced by High-Performance Flash Enclosures (HPFE) Gen2.

Like its predecessors, the DS8900F supports advanced disaster recovery (DR) solutions, business continuity solutions, and thin provisioning.

The IBM DS8910F Rack-Mounted model 993 is described in *IBM DS8910F Model 993 Rack-Mounted Storage System Release 9.1*, REDP-5566.

Note: IBM has recently rebranded its storage portfolio. For more information, see [Evolving the IBM Storage Portfolio Brand Identity and Strategy](#). With this rebranding, IBM DS8900F becomes IBM Storage DS8900F. In this document, IBM Storage DS8900F is also referred to as IBM DS8900F or simply DS8900F.

Authors

This book was produced by a team of specialists from around the world.



Peter Kimmel is an IT Specialist, IBM Redbooks® Project Leader, and the Advanced Technical Skills team lead of the Enterprise Storage Solutions team at the IBM EMEA Storage Competence Center (ESCC) in Frankfurt, Germany. He joined IBM Storage in 1999, and since then has worked with all DS8000 generations, with a focus on architecture and performance. Peter co-authored several DS8000 IBM publications. He holds a Diploma (MSc) degree in physics from the University of Kaiserslautern.



Daniel Beukers has worked in IT for over 20 years, managing and deploying infrastructure. He joined IBM in 2013 to provide storage area network (SAN) and storage support for multiple customers. More recently, he moved into an architect role at Kyndryl, working on the roll-out of a private cloud into new data centers.



Jeff Cook is a DS8000 Subject Matter Expert (SME), and leads the Tucson, Arizona DS8000 Product Engineering team. He has been with IBM for 43 years, providing implementation and technical support to customers and service representatives in complex enterprise environments. For the past 24 years, he has specialized in IBM direct access storage device products, specifically the DS8000 and former enterprise storage systems.



Bozhidar Feraliev holds the role of a DS8000 SME and joined the DS8000 team in 2020. Previously, he was part of the DS8000 Remote Support Team (formerly known as the DS8000 Level-1 support), which he joined in 2016. Bozhidar has experience working closely with clients to troubleshoot any hardware or software issues with a DS8000 product. Before joining IBM, Bozhidar was a technical support engineer for one of the leading internet service providers in Bulgaria.



Jörg Klemm is a Senior Mainframe Consultant working for IBM Platinum Business Partner SVA System Vertrieb Alexander GmbH in Germany. He has over 20 years of experience in IBM working directly with customers. He is primarily focused on enterprise storage, and specializes in Business Continuity solutions. His areas of expertise include Copy Services and IBM Geographically Dispersed Parallel Sysplex® (IBM GDPS®). Jörg has been delivering GDPS solutions for over 15 years.



Connie Riggins is a DS8000 Copy Services and Copy Services Manager SME with the DS8000 Product Engineering group. She started working at IBM® in 2015. Before joining IBM, starting in 1991, Connie worked at Amdahl Corp as a Systems Engineer, and later at Softek Storage Solutions as Product Manager of IBM Transparent Data Migration Facility (TDMF) for IBM z/OS®.



Gaurav Sabharwal joined IBM in 2009. He has 18 years of experience in high-end file and block storage with different vendor products. He works at IBM LBS in India. His areas of expertise include performance analysis; establishing high availability and disaster recovery (HADR) solutions; complex data migration; and implementing storage systems. Gaurav holds a degree in information technology. He also acquired a Post-Graduate Program in Artificial Intelligence and Machine Learning from Texas McCombs Business School.

Thanks to the following people for their contributions to this project:

Alexander Warmuth, Michael Frankenberg, Vasfi Gucer, Jörg Zahn, Marcela Adan
IBM

Thanks to the authors of the previous edition of this book:

Bertrand Dufasne, Sherri Brunson, Lisa Martinez, Mike Stenson
IBM

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:
<https://redbooks.ibm.com/residencies>

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:
redbooks.ibm.com/contact
- ▶ Send your comments in an email to:
redbooks@us.ibm.com

- ▶ Mail your comments to:
IBM Corporation, IBM Redbooks
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:
<http://www.linkedin.com/groups/2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>



Part 1

Concepts and architecture

This part of the book provides an overview of IBM DS8900 concepts and architecture.

This part contains the following chapters:

- ▶ Chapter 1, “Introducing the IBM DS8900F storage system” on page 3
- ▶ Chapter 2, “IBM DS8900F hardware components and architecture” on page 25
- ▶ Chapter 3, “IBM DS8900F reliability, availability, and serviceability” on page 71
- ▶ Chapter 4, “Virtualization concepts” on page 107



Introducing the IBM DS8900F storage system

This chapter introduces the features and functions of the IBM DS8900F storage system that are available with IBM DS8000 Licensed Machine Code (LMC) 7.9.30 (bundle version 89.30.xx.x), which is referred to as Release 9.3. More information about functions and features is provided in subsequent chapters.

This chapter covers the following topics:

- ▶ Introduction to the DS8900F system
- ▶ DS8900F controller options and frames
- ▶ DS8900F architecture and functions overview
- ▶ Performance features

1.1 Introduction to the DS8900F system

The DS8000 family is a high-performance, high-capacity, highly secure, and resilient series of disk storage systems. The DS8900F family is the most advanced of the DS8000 offerings to date. All the DS8900F models are all-flash systems. The high availability (HA) multiplatform support, including IBM Z, and simplified management tools provide a cost-effective path to an on-demand world.

The IBM DS8900 family consists of the following high-performance models:

- ▶ IBM DS8980F: Analytic Class
- ▶ IBM DS8950F: Agility Class
- ▶ IBM DS8910F: Flexibility Class, rack-mounted
- ▶ IBM DS8910F: Flexibility Class, racked

The DS8900F models support the most demanding business applications with their exceptional all-around performance and data throughput. Some models are shown in Figure 1-1.



Figure 1-1 DS8900F frame (standard doors with side covers)

The DS8000 offers features that clients expect from a high-end storage system:

- ▶ High performance
- ▶ High capacity
- ▶ HA
- ▶ Security
- ▶ Cost efficiency
- ▶ Energy efficiency
- ▶ Scalability
- ▶ Business continuity and data protection functions

The DS8900F architecture is server-based. Two powerful POWER9 processor-based servers manage the cache to streamline drive I/Os, which maximizes performance and throughput. These capabilities are further enhanced with the availability of High-Performance Flash Enclosures (HPFEs). HPFE Gen2 is described in *IBM DS8000 High-Performance Flash Enclosure Gen2 (DS8000 R9.0)*, REDP-5422.

The DS8900F is an all-flash storage system that is equipped with encryption-capable flash drives. High-density storage enclosures offer a considerable reduction in the footprint and energy consumption. Figure 1-2 shows the front view of a fully configured DS8950F frame.



Figure 1-2 DS8950F all-flash system with KVM console open

Reduced footprint: The DS8900F is housed in a 19-inch wide rack, and it has a smaller depth than its predecessor, which is the DS8880.

The DS8900F includes a power design that is based on intelligent Power Distribution Units (iPDUs), and non-volatile dual inline memory modules (NVDIMMs) to store data in case of a power outage. The iPDUs allow the DS8900F to achieve the highest energy efficiency in the DS8000 series. The DS8900F complies with ENERGY STAR specifications. All models are available for both three-phase and single-phase power attachment.

I/O enclosures are attached to the POWER9 processor-based servers with Peripheral Component Interconnect® Express (PCIe) Generation 3 cables. The I/O enclosure has six PCIe adapter slots, and two zHyperLink ports.

The rack models also have an integrated keyboard and display that can be accessed from the front of the rack. A pair of small form-factor (SFF) Hardware Management Consoles (HMCs) are installed in the base rack management enclosure. The height of the rack is 40U for all units.

The DS8910F Rack-Mounted model comes without its own rack. It can attach to both mainframe or to distributed systems. When placed in a customer-supplied rack, a keyboard and display can be ordered. When integrated into the IBM 3907 z15 T02 or LinuxONE LR1 business-class mainframe models, the DS8910F Rack-Mounted model shares the rack and the management keyboard and display.

Note: The DS8910F and its specific features are described in the *IBM Power Systems Enterprise AI Solutions*, REDP-5556.

Figure 1-3 on page 7 shows the various components within the base frame. The DS8950F and DS8980F expansion frame enables clients to add more capacity to their storage systems in the same footprint.




Figure 1-3 DS8950F components (base frame)

1.1.1 At-a-glance features of the DS8900F

The DS8900F offers the following features (more details can be found in subsequent chapters of this book):

- ▶ The DS8900F is available with various POWER9 processor options, from a DS8980F with a 22-core per central processor complex (CPC) option to the DS8950F with a 20-core per CPC option to the DS8910F with an 8-core per CPC option. Each system covers a wide range of performance and cost options.
- ▶ The IBM POWER® simultaneous multithreading (SMT) technology allows analytic data processing clients to achieve over 2 million I/O operations per second (IOPS) in database for open environments (70% read and 30% write, 4 KB I/Os, and 50% read cache hit at record-low latencies).
- ▶ Memory configurations of 192 GB - 4.3 TB are available.
- ▶ The DS8900F uses HPFE Gen2 storage enclosures. One HPFE pair can hold 16, 32, or 48 flash drives, with capacities of 800 GB, 1.6 TB, 3.2 TB, 1.92 TB, 3.84 TB, 7.68 TB, and 15.36 TB. The latter four are more cost-efficient versions of high-capacity flash.

- ▶ IBM Z Synergy Services (zsS) include IBM z/OS licensed features that are supported on the storage system. The licensed features include Fibre Channel connection (IBM FICON) attachment for IBM Z, parallel access volume (PAV), HyperPAV, SuperPAV, High-Performance FICON for IBM Z (zHPF), zHyperLink, Transparent Cloud Tiering (TCT), z/OS Distributed Data Backup (zDDB), and an option for thin-provisioned Count Key Data (CKD) volumes.
- ▶ The DS8900F offers enhanced connectivity with 4-port 32 Gbps and 16 Gbps Fibre Channel Protocol (FCP) / FICON host adapters. Each host adapter port can be configured independently for either FCP or FICON protocol.
- ▶ zHPF is an enhancement to the FICON architecture to offload I/O management processing from the IBM Z channel subsystem to the DS8000 host adapters and CPCs. zHPF is an optional feature of IBM Z and DS8000. 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 Db2 I/O is now zHPF-capable and supports the Db2 castout accelerator function that allows the DS8000 to treat a castout as a single chain of I/Os. zHPF improvements like Extended Distance II further accelerate support for heavy write I/Os over an extended distance of up to 100 km (62 miles).
- ▶ DS8900F host adapters include support for IBM Fibre Channel Endpoint Security as part of the cybersecurity solutions that are offered by IBM. Traffic to the host can be authenticated and encrypted. To take advantage of this capability, the platforms that are attached to the DS8900F must also have support for IBM Fibre Channel Endpoint Security. IBM z15 is an example. IBM Fibre Channel Endpoint Security is described in *IBM Fibre Channel Endpoint Security for IBM DS8900F and IBM Z*, SG24-8455.
- ▶ TCT enables direct data movement from the DS8000 CPCs into IBM Cloud® Object Storage, which helps reduce backup workload on IBM Z. Data can be encrypted before leaving the DS8900. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

TCT offers seamless hybrid multicloud integration and can be used to move DS8000 data to an IBM TS7700 Virtualization Engine that is configured as an object store. In this case, the data can be encrypted while it is being transferred, and can be compressed. While configured as object storage, the TS7700 also can still be used with traditional FICON logical tape volumes. For more information, see *IBM DS8000 Transparent Cloud Tiering (DS8000 Release 9.2)*, SG24-8381.
- ▶ DS8900F includes PCIe-based I/O enclosures. Each I/O enclosure connects to both internal servers over a pair of x8 PCIe Gen3 cables, each of which provides 8 Gbps connectivity. Each I/O enclosure also provides two zHyperLink ports.

Each pair of I/O enclosures supports up to two pairs of HPFES. The storage virtualization that is offered by the DS8900F allows organizations to allocate system resources more effectively and better control application quality of service (QoS).
- ▶ The Adaptive Multi-stream Prefetching (AMP) caching algorithm can dramatically improve sequential performance by reducing times for backup, processing for business intelligence, and streaming media. Sequential Adaptive Replacement Cache (SARC) is a caching algorithm that allows you to run different workloads, such as sequential and random workloads, without negatively affecting each other. For example, a sequential workload does not fill up the cache and does not affect cache hits for random workloads. Intelligent Write Caching (IWC) improves the Cache Algorithm for random writes.

- ▶ DS8900F provides a graphical management interface to configure, query status information, and monitor the performance of the DS8000. The DS Storage Management GUI (DS GUI) also provides link buttons for accessing the Service Management GUI and Copy Services Manager (CSM) GUI. Additionally, many service functions that once required using the Service Management GUI can now be accessed from the DS GUI.
- ▶ IBM Easy Tier® is a no-charge feature that enables automatic dynamic data relocation capabilities. Easy Tier optimizes the usage of the flash storage. No manual tuning is required. The auto-balancing algorithms also provide benefits when homogeneous storage pools are used to eliminate hot spots on arrays of drives.
- ▶ Large volume support that allows a DS8000 to support logical unit number (LUN) sizes up to 16 TB. This configuration simplifies storage management tasks. In a z/OS environment, extended address volumes (EAVs) with sizes up to 1 TB are supported.
- ▶ The DS8900F has an Active Volume Protection feature that prevents the deletion of volumes that are still in use.
- ▶ The American National Standards Institute (ANSI) T10 Data Integrity Field (DIF) Protection Information standard is supported by DS8000 for Small Computer System Interface (SCSI) end-to-end data protection on Fixed-Block (FB) architecture LUN volumes for operating systems (OSs) that can use DIF.
- ▶ Dynamic Volume Expansion (DVE) simplifies management by enabling easier, online volume expansion to support application data growth, and to support data center migration and consolidation to larger volumes to ease addressing constraints.
- ▶ Thin provisioning allows the creation of over-provisioned devices for more efficient usage of the storage capacity for open systems. For more information, see *IBM DS8880 Thin Provisioning (Updated for Release 8.5)*, REDP-5343.
- ▶ Quick Initialization provides fast volume initialization for open system LUNs and CKD volumes. Quick Initialization provides access to volumes the moment that the command completes.
- ▶ Full Disk Encryption (FDE), also known as *data at rest encryption*, can protect business-sensitive data by providing drive-based hardware encryption that is combined with external key management software (IBM Security® Guardium® Key Lifecycle Manager, Gemalto SafeNet KeySecure, or Thales Vormetric Data Security Manager). FDE is available for flash drive types. Because encryption is performed by the drive, it is transparent to host systems, and it can be used in any environment, including z/OS.
- ▶ DS8900F R9.2 and later supports a Local Key Management option for FDE on the storage system.
- ▶ DS8900F enables clients to comply with SP-800-131a, which is a US National Institute of Standards and Technology (NIST) directive that provides guidance for protecting sensitive data by using cryptographic algorithms that have key strengths of 112 bits. NIST SP-800-131a mode is configured by default in DS8900F.
- ▶ Payment Card Industry Data Security Standard (PCI DSS) is a feature of the encryption key management process that helps address payment industry requirements.
- ▶ The DS8000 series is certified as meeting the requirements of the IPv6 Ready Logo program. This certification indicates its implementation of IPv6 mandatory core protocols and the ability to interoperate with other IPv6 implementations. The DS8000 can be configured in native IPv6 environments. The logo program provides conformance and interoperability test specifications that are based on open standards to support IPv6 deployment globally. Furthermore, US NIST tested IPv6 with the DS8000 and granted it support for the USGv6 profile and testing program.

- ▶ Lightweight Directory Access Protocol (LDAP) authentication support can simplify user management by allowing the DS8000 to rely on a centralized LDAP directory rather than a local user repository. DS8900F supports native LDAP configuration, and it does not require the preinstalled IBM CSM proxy. The CSM proxy can still be used if you want or to facilitate LDAP for older DS8000 systems. For more information, see *LDAP Authentication for IBM Storage DS8000 Systems: Updated for DS8000 Release 9.3.2*, REDP-5460.
- ▶ For data protection and availability needs, the DS8900F provides a rich set of point-in-time copy (PTC), Remote Mirror, and Remote Copy functions. These functions provide storage mirroring and copying over large distances for disaster recovery (DR) or HA purposes.
- ▶ Safeguarded Copy delivers the ability to create and retain hundreds of PTCs to protect against logical data corruption or malicious destruction. Those copies can be used to verify customer data, analyze the nature of the corruption, and restore critical customer data. Dynamic expansion of Safeguarded volume capacity is now possible. For more information, see *IBM Storage DS8000 Safeguarded Copy (Updated for DS8000 R9.2.3)*, REDP-5506.
- ▶ Support for IBM i variable LUN adds flexibility for volume sizes. You can increase capacity usage for IBM i environments. IBM i volumes can be dynamically expanded.
- ▶ The DS8900F features Smart Rebuild, a function that is designed to reduce the possibility of secondary failures and data loss in RAID arrays. When a predicted failure occurs on one member of a RAID 5 or RAID 6 array, the affected drive is cloned to a spare drive, allowing it to participate in its own rebuild. The cloning process reduces the duration of the rebuild time. The process falls back to traditional rebuild when Smart Rebuild cannot complete.
- ▶ The DS8900F offers enhanced IBM Z Synergy items:
 - Forward Error Correction (FEC)

FEC is a protocol that detects and corrects bit errors that are generated during data transmission. Both the IBM z13® (and later) and the DS8000 extend the usage of FEC to complete end-to-end coverage for 16 Gbps and faster links and preserve data integrity with more redundancy.
 - Fibre Channel (FC) Read Diagnostic Parameters (RDP)

FC RDP improves the end-to-end link fault isolation for 16 Gbps and faster links on current IBM Z and DS8000 systems. RDP data provides the optical signal strength, error counters, and other critical information that is crucial to determine the quality of the link.
 - FICON® Dynamic Routing (FIDR)

FIDR is an IBM Z and DS8000 feature that supports the use of dynamic routing policies in the fabric switch to balance load across inter-switch links (ISLs) on a per I/O basis.
 - Fabric I/O Priority

Fabric I/O Priority is an end-to-end synergy feature among the z/OS Workload Manager (WLM), Brocade storage area network (SAN) Fabric, and the storage system to manage QoS on a single I/O level.
 - IBM Fibre Channel Endpoint Security

As one of the first hosts, and following a pervasive encryption philosophy, IBM z15® supports SAN authentication and FC encryption of FICON traffic to the DS8900F.
 - IBM zHyperLink

zHyperLink is a short distance link technology that dramatically reduces latency by interconnecting the IBM Z CPCs directly to the I/O bays of the DS8900.

All DS8900F models can support zHyperLink. The current zHyperLink release supports read/write I/O. zHyperLink is intended to complement FICON technology to accelerate I/O requests that are typically used for transaction processing.

For more information, see *Getting Started with IBM zHyperLink for z/OS*, REDP-5493.

Note: For more information about the IBM Z synergy features, see *IBM DS8900F and IBM Z Synergy DS8900F: Release 9.3 and z/OS 2.5*, REDP-5186.

1.2 DS8900F controller options and frames

The DS8900F family consists of a series of distinct models that are all flash storage:

- ▶ DS8980F machine type 5341 model 998 and DS8980F model E96
- ▶ DS8950F machine type 5341 model 996 and DS8950F model E96
- ▶ DS8910F machine type 5341 model 994 as racked in a single frame
- ▶ DS8910F machine type 5341 model 993 as the rack-mounted model

The DS8900F includes the following features:

- ▶ IBM POWER9 processor-based server technology

The DS8900F features the IBM POWER9 processor-based server technology for high performance. Compared to the IBM POWER8® processors that were used in the DS8880 servers, IBM POWER9 processor-based servers excel with a smaller front-end latency and up to 60% more IOPS in transaction-processing workload environments. The number of transistors per core has almost doubled between these two POWER generations.

- ▶ Nondisruptive upgrade path

A nondisruptive upgrade path for the DS8900F allows processor, cache, host adapters, storage upgrades, and expansion frame upgrades to be installed concurrently without disrupting applications.

- ▶ The air-flow system allows optimal horizontal cool down of the storage system. The DS8900F is designed for hot- and cold-aisle data center design, drawing air for cooling from the front of the system and exhausting hot air at the rear. For more information, see 2.6, “Power and cooling” on page 58.

- ▶ The DS8900F offers three system class options:

- DS8980F Mod 998 - Analytic Class, which is based on the IBM Power S924 architecture.
- DS8950F Mod 996 - Agility Class, which is based on the IBM Power S924 architecture.
- DS8810F Mod 994 and Mod 993 - Flexibility Class, which is based on IBM Power S922 architecture.

For more information about the specifics for each model, see 2.3.1, “IBM POWER9 processor-based CPCs” on page 42.

- ▶ HPFE

HPFE Gen2 is an enclosure that can support both high-performance flash drives with capacities of 800 GB - 3.2 TB and high-capacity flash drives of 1.92 TB, 3.84 TB, 7.68 TB, or 15.36 TB (2.5-inch form-factor) in a 2U rack space. The enclosures must be installed in pairs.

For more information, see *IBM DS8000 High-Performance Flash Enclosure Gen2 (DS8000 R9.0)*, REDP-5422.

- ▶ Rack-mountable and racked options

The DS8900F comes in racks with a footprint that is reduced compared its DS8880 predecessor models. To use the DS8900F in a customer-provided rack, use the DS8910F model 993.

1.3 DS8900F architecture and functions overview

The DS8900F offers continuity with the fundamental architecture of its predecessors: the DS8880 and DS8870 models. This architecture ensures that the DS8900F can use a stable and proven operating environment that offers optimal availability. The hardware is optimized to provide higher performance, connectivity, and reliability.

1.3.1 Overall architecture and components

For more information about the available configurations for the DS8900F, see Chapter 2, “IBM DS8900F hardware components and architecture” on page 25.

Note: Some technical aspects are specific to the DS8910F rack-mounted model. For a full overview of the architectural aspects of the DS8910F rack-mounted model, see *IBM DS8910F Model 993 Rack-Mounted Storage System Release 9.1*, REDP-5566.

IBM POWER9 processor technology

The POWER9 processor is manufactured by using the IBM 14 nm Silicon-On-Insulator (SOI) technology. Each chip is 693 mm² and contains 8 billion transistors. The DS8900F uses the PCIe I/O controllers and an interconnection system that connects all components within the chip. POWER9 processor advancements in multi-core and multithreading are remarkable. These multithreading capabilities improve the I/O throughput of the DS8900F storage servers.

The DS8900F family offers several configurations of CPCs:

- ▶ The DS8980F model has two 22-core CPCs (each CPC has two 11-core processors) with a total of 4.3 TB of system memory.
- ▶ The DS8950F offers two configurations:
 - 10-core per CPC with a total of 512 GB of total system memory
 - 20-core per CPC with a total of 1, 2, or 3.4 TB of total system memory
- ▶ The DS8910F models have an 8-core processor configuration per CPC with 192 GB or 512 GB of system memory.

A CPC is also referred to as a *storage server*. For more information, see Chapter 3, “IBM DS8900F reliability, availability, and serviceability” on page 71.

Internal PCIe-based fabric

The DS8900F fabric includes the following specifications:

- ▶ DS8900F POWER9 processor-based servers are based on the current PCIe Gen4 architecture to provide up to 16-lane (x16) high-speed connectivity to internal adapters.
- ▶ PCIe adapters provide point-to-point connectivity to the I/O enclosures. The I/O enclosures provide connectivity between the I/O adapters and the POWER9 processor complexes.
- ▶ The I/O enclosures provide PCIe Gen3 connectivity to all installed host and device adapters. Each I/O enclosure features six PCIe x8 adapter slots and two extra PCIe x8 connectors for attachment to zHyperLink optical transceivers.

For more information, see Chapter 2, “IBM DS8900F hardware components and architecture” on page 25.

HPFE Gen2

The HPFE flash RAID adapters are installed in pairs and split across an I/O enclosure pair. They occupy the third and sixth PCIe slots according to the adapter pair plugging order.

HPFE drive enclosures are also installed in pairs, and connected to the corresponding flash RAID adapter pair over eight 6 Gbps SAS cables for high bandwidth and redundancy. Each drive enclosure can contain up to twenty-four 2.5-inch SAS flash drives. Flash drives are installed in groups of 16, and split evenly across the two drive enclosures in the pair.

Each flash adapter pair and HPFE pair deliver up to 900 K IOPS reads, 225 K IOPS writes, and up to 14 GBps (read) and 9.5 GBps (write) bandwidth.

For more information, see Chapter 2, “IBM DS8900F hardware components and architecture” on page 25.

Drive options

Flash drives provide up to 100 times the throughput and 10 times lower response time than 15 K revolutions per minute (RPM) hard disk drives (HDDs). They also use less power than traditional HDDs. For more information, see Chapter 5, “IBM DS8900F physical planning and installation” on page 141.

Flash drives are grouped into three tiers, based on performance and capacity. These flash Drives are supported across all DS8900F models.

- ▶ 2.5-inch flash Tier 0 high-performance flash drives:
 - 800 GB
 - 1.6 TB
 - 3.2 TB
- ▶ 2.5-inch flash Tier 1 high-capacity flash drives:
 - 3.84 TB
- ▶ 2.5-inch flash Tier 2 high-capacity flash drives:
 - 1.92 TB
 - 7.68 TB
 - 15.36 TB

All flash drives in the DS8900F are encryption-capable. Enabling encryption is optional, and requires at least two external key servers or the local key management feature.

Easy Tier

Easy Tier enables the DS8000 to automatically balance data placement on disk drives to avoid hot spots on flash arrays. Easy Tier can place data in the storage tier that best suits the access frequency of the data. Highly accessed data can be moved nondisruptively to a higher tier, for example, to 1.6 TB Flash Tier 0 drives. Cold data or data that is primarily accessed sequentially is moved to a lower tier of high-capacity drives. Easy Tier includes more components:

- ▶ *Easy Tier Application* is an application-aware storage utility to help deploy storage more efficiently by enabling applications and middleware to direct more optimal placement of the data by communicating important information about current workload activity and application performance requirements. It is possible for Db2 applications in z/OS environments to give hints of data placement to Easy Tier at the data set level.
- ▶ *Easy Tier Heat Map Transfer (HMT)* can take the data placement algorithm on the Metro Mirror (MM) Global Copy (GC) and Global Mirror (GM) primary site and reapply it to the MM, GC, or GM secondary site when failover occurs by using the Easy Tier Heat Map Transfer Utility (HMTU). With this capability, the DS8000 systems can maintain application-level performance. The Easy Tier HMT functions support Metro/Global Mirror (MGM) to transfer a heat map automatically to a tertiary site.

Note: Easy Tier Server was removed from marketing support. It was replaced by the Flash Cache option of IBM AIX® 7.2.

The following Easy Tier controls are also available:

- ▶ For Easy Tier to function effectively, it needs some free extents in each extent pool to be able to move around extents. Rather than having the storage administrator monitor capacity, the system can reserve some space for Easy Tier extent movements. You can enable space reservation for Easy Tier by running `chsi -etsrmode enable` (enabled by default).
- ▶ Another control is the allocation policy for new volumes. You can control the allocation policy according to your needs by running the `chsi` command with `-ettierorder highutil` or `highperf`. The data allocation order for all flash systems can be changed between the following settings:
 - High-Utilization: Flash Tier 1 → Flash Tier 2 → Flash Tier 0
 - High-Performance: Flash Tier 0 → Flash Tier 1 → Flash Tier 2 (the default)

For more information about Easy Tier features, see the following resources:

- ▶ *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667
- ▶ *DS8870 Easy Tier Application*, REDP-5014
- ▶ *IBM DS8870 Easy Tier Heat Map Transfer*, REDP-5015

Host adapters

The DS8900F offers 32 Gbps and 16 Gbps host adapters. Both types have four ports each, and each port can be independently configured for either FCP or FICON:

- ▶ The 32 Gbps adapter has four ports. Each port independently auto-negotiates to an 8, 16, or 32 Gbps link speed.
- ▶ The 16 Gbps adapter has four ports. Each port independently auto-negotiates to a 4, 8, or 16 Gbps link speed.

For more information, see Chapter 2, “IBM DS8900F hardware components and architecture” on page 25.

Storage Hardware Management Console for the DS8000

HMCs are the focal point for notification, management, and maintenance activities. HMCs proactively monitor the state of your system and notify you and IBM when service is required. HMCs can also be used for management of copy services by using the preinstalled IBM Copy Services Manager (CSM) software.

Every DS8900F includes two HMCs for redundancy, which are installed in the management enclosure in the base rack. DS8900F HMCs support IPv4 and IPv6 standards. For more information, see Chapter 6, “IBM DS8900F Management Console planning and setup” on page 167.

Isolated key server

The DS8900F includes FDE flash drives. To configure a DS8900F to use data at rest encryption, at least two key servers are required. An isolated key server with dedicated hardware and non-encrypted storage resources is required. It can be ordered from IBM. For more information, see 5.3.6, “Key manager servers for encryption” on page 161.

You can also encrypt data before it is transmitted to the cloud when using the TCT feature. For more information, see the *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security (DS8000 Release 9.2)*, REDP-4500.

Local Key Manager

Release 9.2 and 9.3 offer Local Key Management for encryption. Local Key Management provides a DS8000 encryption and key management solution to minimize the risk of exposure and reduce operational costs. Local key encryption offers good security for data on disk, even when a drive is physically removed with the intent of stealing the data. For more information about local key manager, see *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security (DS8000 Release 9.2)*, REDP-4500.

1.3.2 Storage capacity

The physical storage capacity for the DS8900F is installed in fixed increments that are called *drive sets* or *flash drive sets*. A drive set contains 16 flash drives, all of which have the same capacity and the same drive class. Both high-performance and high-capacity flash drives are available in sets of 16.

The available drive options provide industry-class capacity and performance to address a wide range of business requirements. The DS8000 storage arrays can be configured as RAID 6, RAID 10, or RAID 5, depending on the drive type.

RAID 6 is now the default and preferred setting for the DS8900F. RAID 5 can be configured for drives of less than 1 TB, but this configuration is not preferred and requires a risk acceptance, and a field Request for Price Quotation (RPQ) for enterprise hard disk drive (HDD) drives. Flash Tier 0 drive sizes larger than 1 TB can be configured by using RAID 5, but require an RPQ and an internal control switch to be enabled. RAID 10 continues to be an option for all-flash drives.

For more information, see 2.2, “DS8900F configurations and models” on page 28.

1.3.3 Supported environments

The DS8000 offers connectivity support across a broad range of server environments, including IBM Power, IBM Z, servers from HPE and Oracle, AMD-based, and Intel-based x64 servers.

The DS8000 supports over 60 platforms. For the list of supported platforms, see the [IBM System Storage Interoperation Center \(SSIC\) for DS8000](#).

This rich support of heterogeneous environments and attachments, with the flexibility to partition easily the DS8000 storage capacity among the attached environments, can help support storage consolidation requirements and dynamic environments.

1.3.4 Configuration flexibility

The DS8000 series uses virtualization techniques to separate the logical view of hosts onto LUNs from the underlying physical layer, providing high configuration flexibility. For more information about virtualization, see Chapter 4, “Virtualization concepts” on page 107.

Dynamic Volume Expansion

DVE increases the capacity of open systems, IBM i, and IBM Z volumes while the volumes remain connected to a host system. This capability simplifies data growth by providing volume expansion without taking volumes offline. Certain OSs do not support a change in volume size. Therefore, a host action is required to detect the change after the volume capacity is increased.

Large LUN and large Count Key Data volume support

You can configure LUNs and volumes to span arrays, allowing for larger LUN sizes of up to 16 TB in open systems.

Tip: Copy Services (CS) are currently supported for LUN sizes of up to 4 TB.

The maximum CKD volume size is 1,182,006 cylinders (1 TB), which can greatly reduce the number of volumes that are managed. This large CKD volume type is called a 3390 Model A. It is referred to as an Extended Address Volume (EAV).

T10 Data Integrity Field support

The DS8900F supports the T10 DIF standard for FB volumes that are accessed by the FCP channel of Linux on IBM Z and IBM AIX on IBM Power. You can define LUNs with an option to instruct the DS8000 to use the CRC-16 T10 DIF algorithm to store the data. You can also create T10 DIF-capable LUNs. The support for IBM i variable LUNs adds flexibility for volume sizes and can increase capacity usage for IBM i environments.

VMware vStorage API for Array Integration support

The VMware vStorage API for Array Integration (VAAI) feature offloads specific storage operations to the storage system for greatly improved performance and efficiency. With VAAI, VMware vSphere can perform key operations faster and use less CPU, memory, and storage bandwidth.

The DS8000 supports the following VAAI primitives:

- ▶ Atomic Test and Set (ATS), which is also known as *Compare and Write* for hardware-assisted locking.
- ▶ Clone Blocks (Extended Copy or *XCOPY*) for hardware-assisted move or cloning. For XCOPY, the DS8000 CS license is required. Also, XCOPY is not supported by extent space efficient (ESE) volumes and volumes larger than 4 TB; the target of an XCOPY cannot be an MM or GM primary volume.

- ▶ Write Same (or *Block Zero*) is used to initialize new volumes and is supported by all volumes.
- ▶ Space Release (*UNMAP*) is supported only with ESE volumes that use small extents (16 mebibytes (MiB)).

IBM DS8000 Storage Replication Adapter

IBM DS8000 Storage Replication Adapter (SRA) is a software add-on that integrates with the VMware vCenter Site Recovery Manager (SRM) solution and enables SRM to perform failovers together with DS8000 storage systems. The DS8000 SRA extends SRM capabilities and allows it to employ DS8000 replication and mirroring as part of the SRM comprehensive disaster recovery planning (DRP) solution.

OpenStack

The DS8000 supports the OpenStack cloud management software for business-critical private, hybrid, and public cloud deployments. The DS8900F supports features in the OpenStack environment, such as volume replication and volume retype. The Cinder driver for DS8000 is now open source in the OpenStack community. The `/etc/cinder.conf` file can be directly edited for the DS8000 back-end information.

For more information about the DS8000 and OpenStack, see *Using IBM DS8000 in an OpenStack Environment*, REDP-5220.

Red Hat OpenShift

Red Hat OpenShift is an open source container application platform that is based on the Kubernetes container orchestrator for enterprise application development and deployment. IBM Red Hat OpenShift Container Platform (OCP) provides developers and IT organizations with a hybrid cloud application platform for deploying new and existing applications on secure, scalable resources with minimal configuration and management overhead. OCP supports various programming languages and frameworks, such as Java, JavaScript, Python, Ruby, and PHP.

DS8900F supports the Container Storage Interface (CSI) specification. IBM released an open source CSI driver for IBM storage that allows dynamic provisioning of storage volumes for containers on Kubernetes and IBM Red Hat OpenShift Container Platform (OCP).

The CSI driver for IBM block storage systems enables container orchestrators such as Kubernetes to manage the lifecycle of persistent storage. This CSI is the official operator to deploy and manage the IBM block storage CSI driver.

For more information about CSI, see [IBM Documentation](#).

RESTful application programming interface

With the DS8000 support of RESTful application programming interface (API) services, DS8900F clients or cloud administrators can design and implement the DS8000 management applications by using the Representational State Transfer (REST) software architecture.

For more information about the RESTful API, see *Exploring the DS8870 RESTful API Implementation*, REDP-5187.

Flexible LUN-to-LSS association

With no predefined association of arrays to logical subsystems (LSSs) on the DS8000 series, users can put LUNs or CKD volumes into LSSs and better use the 256 address range, particularly for IBM Z.

Simplified LUN masking

In the new GUI, LUNs are directly mapped to the host, and the user cannot define volume groups. Volume groups still exist on the DS8900F systems, but they are not visible from the GUI because they are created in the background during the assignment of a volume to the host.

Thin-provisioning features

Volumes in the DS8900F can be provisioned as full or thin. When clients plan capacity, they must consider the number of volumes in the extent pool (or overall storage system) and the degree of over-allocation that is planned for.

These volumes feature enabled over-provisioning capabilities that provide more efficient usage of the storage capacity and reduced storage management requirements. For more information, see Chapter 4, “Virtualization concepts” on page 107 and *IBM DS8880 Thin Provisioning (Updated for Release 8.5)*, REDP-5343.

Maximum values of logical definitions

The DS8000 features the following maximum values for the major logical definitions:

- ▶ Up to 255 LSSs
- ▶ Up to 65,280 logical devices
- ▶ Up to 16 TiB LUNs
- ▶ Up to 1,182,006 cylinder (1 TB) CKD volumes
- ▶ Up to 130,560 FICON logical paths (512 logical paths for each control unit image) on the DS8000
- ▶ Up to 1,280 logical paths for each FC port
- ▶ Up to 8,192 process logins (509 for each SCSI-FCP port)

1.3.5 Copy Services functions

For IT environments that cannot afford to stop their systems for backups, the DS8000 provides IBM FlashCopy®. This fast replication technique can provide a PTC of the data in a few seconds or even less. A recent addition is the ability to perform Cascaded FlashCopy backups, where the target of a FlashCopy relationship can also be the source of another FlashCopy relationship. For more information, see *DS8000 Cascading FlashCopy Design and Scenarios*, REDP-5463.

To meet the challenges of cybersecurity, the Safeguarded Copy function, based on the FlashCopy technology, can create and retain hundreds of PTCs for protection against logical data corruption. Release 9.2 added the capability to restore a recovered Safeguarded Copy to a production copy of the data. For more information, see *IBM Storage DS8000 Safeguarded Copy (Updated for DS8000 R9.2.3)*, REDP-5506.

For data protection and availability needs, the DS8900F provides MM, GM, GC, MGM, and z/OS Global Mirror (zGM), which are Remote Mirror and Remote Copy functions. These functions are also available and are fully interoperable with previous models of the DS8000 family. These functions provide storage mirroring and copying over large distances for DR or availability purposes.

CS scope limiting is the ability to specify policy-based limitations on CS requests.

For more information about CS, see *IBM DS8000 Copy Services: Updated for IBM DS8000 Release 9.1*, SG24-8367.

1.3.6 Service and setup

Beginning with Release 9.3, DS8000 systems have a single machine type that is called 5341. The former warranty and service options are now offered as part of Expert Care. Options range from a 1-year base warranty to 5-year Expert Care Advanced or Premium, which may be chosen when you order a new system or upgrade as requirements change.

Figure 1-4 describes the available Expert Care options.

Expert Care Services	Base Warranty 5341	Advanced 5131 (A01-A05)	Premium 5131 (P01-P05)
Hardware Maintenance (IOR = IBM On-site Repair)	1-Year 24x7 same day	24x7 same day	24x7 same day
Support Line (24x7 remote technical support)		✓	✓
Predictive Support		✓	✓
Technical Account Manager (TAM)			✓
Enhanced Response Time (30 min for Severity 1 & 2)			✓
Remote Code Load (up to 2x per year)			✓
Outside the bundle			
On-site Code Load (optional add-on via Feature Code)		Optional add-on	Optional add-on
Media Retention (optional add-on via Service Pac)		Optional add-on	Optional add-on

Figure 1-4 Expert Care: 1-year and 5-year options

Physical installation of the DS8000 is performed by IBM by using the installation procedure for this system. The client's responsibility is the installation planning, retrieval, and installation of feature activation codes, logical configuration, and execution.

The storage system HMC is the focal point for maintenance and service operations. Two HMCs are inside the DS8900F management enclosure, and they continually monitor the state of the system. HMCs notify IBM, and they can be configured to notify you when service is required.

The HMC is also the interface for Call Home and remote support, which can be configured to meet client requirements. It is possible to allow one or more of the following configurations:

- ▶ Call Home on error (machine-detected)
- ▶ Remote connection for a few days (client-initiated)
- ▶ Remote problem log collection (service-initiated)

The remote connection between the HMC and IBM Support is performed by using the Assist On-site (AOS) feature. AOS offers more options, such as Transport Layer Security (TLS), and enhanced audit logging. For more information, see *IBM Assist On-site for Storage Overview*, REDP-4889.

IBM Remote Support Center (RSC) is also an available option for providing IBM Support remote support access to systems.

For Release 9.3 systems, IBM provides three options for microcode updates:

- ▶ Customer Code Load
- ▶ Remote Code Load (RCL)
- ▶ Onsite SSR Code Load

For customers who choose the base warranty service or Expert Care Advanced, Customer Code Load is the default method for performing concurrent microcode updates:

- ▶ Microcode bundles are downloaded and activated by the customer by using the standard DS Storage Manager GUI.
- ▶ The download defaults to the current recommended bundle, or an alternative compatible bundle may be chosen.
- ▶ Health checks are run before the download, and again before activation to ensure that the system is in good health.
- ▶ If a problem is encountered anywhere in the process, a ticket is opened automatically with IBM Support, and the ticket number is provided in the GUI for reference.
- ▶ After the problem is corrected, the code load can be restarted, and automatically resumes after the last successful step.

Customers who want to have an IBM Systems Service Representative (IBM SSR) perform Onsite Code Load may purchase Feature Code #AHY2 with Expert Care Premium, or Feature Code #AHY3 with Expert Care Advanced.

1.3.7 IBM Certified Secure Data Overwrite

IBM Certified Secure Data Overwrite (SDO) is a process that provides a secure overwrite of all data storage in a DS8900F storage system. Before you perform a SDO, you must remove all the logical configuration and any encryption groups that may be configured. The process is then initiated by the IBM SSR. For more information, see 3.7.3, “IBM Certified Secure Data Overwrite” on page 101.

1.4 Performance features

The DS8900F offers optimally balanced performance. The DS8000 incorporates many performance enhancements, such as a dual multi-core IBM POWER9 processor complex implementation, fast 32 Gbps and 16 Gbps FCP / FICON host adapters, HPFE dedicated flash architecture in second generation, classical flash drives, and high-bandwidth, fault-tolerant point-to-point PCIe internal connections.

With all of these components, the DS8900F is positioned at the top of the high-performance category.

1.4.1 4.3 TB cache

DS8900F supports up to a 4.3 TB cache, which can improve the cache hit ratio and result in the reduction of the response time of the same workload demand. For a cache size sensitive workload with an improved cache hit ratio, the larger cache size provides higher throughput with lower response time.

1.4.2 32 Gbps and 16 Gbps host adapters

The DS8900F supports 32 Gbps and 16 Gbps host adapters. This connectivity reduces latency and provides faster single stream and per-port throughput. These adapters can work up to two-speed classes below their own nominal class, but they do not support Fibre Channel Arbitrated Loop (FC-AL) connections. The *Lights on Fastload* feature avoids path disturbance during code loads.

The 16 Gbps host bus adapter (HBA) supports IBM Fibre Channel Endpoint Security authentication. The 32 Gbps HBA supports both IBM Fibre Channel Endpoint Security authentication and line-rate encryption.

1.4.3 Sophisticated caching algorithms

IBM Research® conducts extensive investigations into improved algorithms for cache management and overall system performance improvements. To implement sophisticated caching algorithms, it is essential to include powerful processors for the cache management. With a 4 KB cache segment size and up to 4.3 TB overall cache sizes, the tables to maintain the cache segments become large.

Sequential Prefetching in Adaptive Replacement Cache

One of the performance features of the DS8000 is its self-learning cache algorithm, which optimizes cache efficiency and enhances cache hit ratios. This algorithm, which is used in the DS8000 series, is called *Sequential Prefetching in Adaptive Replacement Cache* (SARC).

SARC provides the following abilities:

- ▶ Sophisticated algorithms to determine the data to store in cache that is based on recent access and the frequency needs of the hosts.
- ▶ Prefetching, which anticipates data before a host request and loads it into cache.
- ▶ Self-learning algorithms to adaptively and dynamically learn the data to store in cache that is based on the frequency needs of the hosts.

Adaptive Multi-stream Prefetching

AMP is a breakthrough caching technology that improves performance for common sequential and batch processing workloads on the DS8000. AMP optimizes cache efficiency by incorporating an autonomic, workload-responsive, and self-optimizing prefetching technology.

Intelligent Write Caching

IWC improves performance through better write-cache management and destaging the order of writes. IWC can also double the throughput for random write workloads. Specifically, database workloads benefit from this new IWC cache algorithm.

SARC, AMP, and IWC play complementary roles. Although SARC carefully divides the cache between the RANDOM and the SEQ lists to maximize the overall hit ratio, AMP is managing the contents of the SEQ list to maximize the throughput that is obtained for the sequential workloads. IWC manages the write cache and decides the order and rate to destage to disk.

1.4.4 Flash storage

Today's installations are mostly hybrid multitier installations, mixing more cost-economical large-capacity drives with high-performance flash together in one system. Given the capacity increase of the flash modules, the price drop of capacity-optimized flash compared to high-RPM HDDs, and given their savings on energy consumption and space, most clients decide on an all-flash array storage strategy.

The DS8900F flash storage can be tiered, with three tiers of flash storage that are available. Then, you can use Easy Tier to optimize the storage. The DS8900F offers automated algorithms that optimize the tiering and place hot areas onto higher-tiered flash arrays.

To improve data transfer rate (IOPS) and response time, the DS8900F supports flash drives and high-performance flash drives, which are based on NAND technology. With the flash drives and the specific architecture that is used in the HPFEs, much higher IOPS densities (IOPS per GB) are possible than with ordinary solid-state drives (SSDs).

Flash drives sharply improve I/O transaction-based performance over traditional HDDs in standard drive enclosures.

High-performance flash drives use the flash RAID adapters in the I/O enclosures, and PCIe connections to the processor complexes. The high-performance flash drive types are high-IOPS class enterprise storage devices that are targeted at flash Tier 0, for I/O-intensive workload applications that need high-level, fast-access storage. The high-capacity flash drive types for flash Tiers 1 and 2 often have an acquisition price point that helps eliminate HDDs when replacing a storage system.

Flash drives offer many potential benefits over HDDs, including better IOPS, lower power consumption, less heat generation, and lower acoustical noise. For more information, see Chapter 5, "IBM DS8900F physical planning and installation" on page 141.

1.4.5 Performance for IBM Z

The DS8000 series supports the following IBM performance enhancements for IBM Z environments:

- ▶ PAVs enable a single IBM Z server to simultaneously process multiple I/O operations to the same logical volume, which can reduce device queue delays. This reduction is achieved by defining multiple addresses for each volume. With Dynamic PAV, the assignment of addresses to volumes can be automatically managed to help the workload meet its performance objectives and reduce overall queuing.
- ▶ HyperPAV enables applications to achieve equal or better performance than with PAV alone while using fewer unit control blocks (UCBs) and eliminating the latency in targeting an alias to a base. With HyperPAV, the system can react immediately to changing I/O workloads.
- ▶ SuperPAV allows z/OS to use an *alias address* from another logical control unit (LCU).
- ▶ Multiple Allegiance expands the simultaneous logical volume access capability across multiple IBM Z servers. This function, with PAV, enables the DS8000 to process more I/Os in parallel, which improves performance and enables greater use of large volumes.
- ▶ I/O priority queuing allows the DS8000 series to use I/O priority information that is provided by the z/OS WLM to manage the processing sequence of I/O operations at the adapter level.
- ▶ zHPF enables multiple channel commands to be sent to the control unit as a single entity. The channel forwards a chain of commands, and it does not need to track each individual channel command word (CCW). This configuration improves FICON throughput on the DS8000 I/O ports. The DS8000 systems also support the new zHPF I/O commands for multi-track I/O operations, Db2® list-prefetch, sequential access methods, and Db2 castout acceleration.
- ▶ zHyperWrite is another enhancement for Db2 clients. In a z/OS MM environment, it enables Db2 log updates to be written to the primary and secondary volumes in parallel. This configuration reduces the latency for log writes, and so improving transactional response times and log throughput. The MM primary volume needs to be enabled with IBM HyperSwap® by either IBM Geographically Dispersed Parallel Sysplex (IBM GDPS) or IBM Copy Services Manager.
- ▶ zHyperLink provides a connectivity method that dramatically reduces latency by interconnecting the IBM Z system directly to the I/O bay of the DS8900. Response times lower than 20 μ s are possible for qualifying I/Os. zHyperLink supports read/write I/O.
- ▶ zHyperLink writes to the GM primary extends the benefits of zHyperLink to GM environments. With 32 Gbps FICON connectivity, where a zHPF write operation to a GM takes about 100 μ s, that same write operation takes only 27 μ s over zHyperLink.
- ▶ Release 9.2 added the ability to allow consistent read from MM secondary volumes over zHyperLink or zHPF in a HyperSwap configuration.
- ▶ With the TCT function, the DS8000 CPCs can be scheduled to directly move data over to Cloud Object Storage or to tape without data going through the host, which can sharply reduce IBM Z MIPS usage during the backup periods and reduce expenses while simplifying data archiving operations. The TCT function now enables multi-cloud support where up to eight cloud targets can be defined for each DS8900F. For more information, see *IBM DS8000 Transparent Cloud Tiering (DS8000 Release 9.2)*, SG24-8381.

For more information about performance on IBM Z, see *IBM DS8900F and IBM Z Synergy DS8900F: Release 9.3 and z/OS 2.5*, REDP-5186.

1.4.6 Performance enhancements for IBM Power servers

Many IBM Power users can benefit from the following DS8000 features:

- ▶ End-to-end I/O priorities
- ▶ Cooperative caching
- ▶ Long busy wait host tolerance
- ▶ Automatic port queues



IBM DS8900F hardware components and architecture

This chapter describes the hardware components of the IBM DS8900F. It provides insights into the architecture and individual components.

This chapter covers the following topics:

- ▶ Flash drive terminology of the DS8900F
- ▶ DS8900F configurations and models
- ▶ DS8900F architecture overview
- ▶ I/O enclosures and adapters
- ▶ Flash drive enclosures
- ▶ Power and cooling
- ▶ Management Console and network

Note: The IBM DS8910F model 993 Rack-Mounted system has some configuration differences from the racked model 994. For more information, see *IBM DS8910F Model 993 Rack-Mounted Storage System Release 9.1*, REDP-5566.

2.1 Flash drive terminology of the DS8900F

This section describes the naming conventions that are used to describe the DS8900F components and features. Although most of these terms are introduced in other chapters of this book, they are repeated and summarized here.

2.1.1 Storage system

The term *storage system* in this context describes a single DS8900F (base frame plus optional expansion frame).

Base frame

The DS8900F has four available base frame models in three DS8900F families (Analytic class, Agility class, and Flexibility Class). The model numbers depend on the hardware configuration for each: DS8980F, DS8950F, and DS8910F. In this chapter, the DS8900F family name, or model number, are used interchangeably. Table 2-1 lists each of the frame models.

Table 2-1 DS8900F frame models and expansion frames

DS8900F	Base frame model	Expansion frame model	Max expansion frames
DS8980F	998	E96	1
DS8950F	996	E96	1
DS8910F	994	N/A	N/A
DS8910F	993	N/A	N/A

Each base frame is equipped with dual Hardware Management Consoles (HMCs). To increase the storage capacity and connectivity, an expansion frame can be added to any DS8980F model 998, or a DS8950F model 996 with 40 cores, and at least 1 TB system memory.

For more information about the base frame configuration, see 2.2.5, “DS8900F base frames” on page 35.

For more information about the DS8910F model 993 Rack-Mounted system, see *IBM DS8910F Model 993 Rack-Mounted Storage System Release 9.1*, REDP-5566.

Expansion frame

The DS8980F and DS8950F support one optional expansion frame, which provides space for extra storage capacity and also supports up to two extra I/O enclosure pairs. To add an expansion frame to the DS8950F, the storage system must first be configured with 1024 GB of memory and 40 processor cores.

With these models, you can place the expansion frame up to 20 meters away from the base frame. To use this feature, use the optical Peripheral Component Interconnect Express (PCIe) I/O Bay interconnect. The Copper PCIe I/O Bay interconnect is used when the expansion frame is physically next to the base frame. For more information about the expansion frame connections, see 2.2.6, “DS8900F expansion frame” on page 38.

All DS8900F system memory and processor upgrades can be performed concurrently.

2.1.2 Management Enclosure

The Management Enclosure (ME) contains two HMC systems that provide internal network and power control communications to the central processor complexes (CPCs), I/O enclosures, and intelligent Power Distribution Units (iPDUs).

A monitor and keyboard with trackpad are provided for local management. The ME also provides up to two Ethernet connections from each HMC for remote management.

Figure 2-1 shows the front view of the 2U ME.



Figure 2-1 DS8900F Management Enclosure

2.1.3 Central processor complex

The DS8900F uses two POWER9-based servers, which are referred to as CPCs. Internal server, processor complex, or central electronics complex (CEC) are also sometimes used. For more information, see 2.3.1, “IBM POWER9 processor-based CPCs” on page 42.

The characteristics for CPCs for each model type are listed in Table 2-2.

Table 2-2 DS8900F processor and cache details

Model	Processors per storage system	Cache per storage system
998	Forty-four cores	4352 GB
996	Forty cores	512 GB
		1024 GB
		2048 GB
		3456 GB
993 and 994	Sixteen cores	192 GB
		512 GB

Both CPCs in a DS8900F system share the system workload. The CPCs are redundant, and either CPC can fail over to the other for scheduled maintenance, upgrade tasks, or if a failure occurs. The CPCs are identified as CPC 0 and CPC 1. A logical partition (LPAR) in each CPC runs the AIX V7.x operating system (OS) and storage-specific Licensed Internal Code (LIC). This LPAR is called the *storage node*. The storage servers are identified as *Node 0* and *Node 1* or *server0* and *server1*.

2.2 DS8900F configurations and models

This section presents the DS8900F configurations and models at the time of writing. The DS8900F consists of one base frame, and for the DS8980F or DS8950F, one optional expansion frame. The CPCs are in the base frame, and they can be upgraded with more processor cores and system memory to accommodate growing performance or when more storage capacity or host connectivity is required.

The main variations between models are the combinations of CPCs, I/O enclosures, storage enclosures, and flash drives. System memory, processors, storage capacity, and host attachment upgrades from the smallest to the largest configuration can be performed concurrently.

Beginning with Release 9.3 new builds, DS8900F storage systems use machine type 5341. The former warranty and service options are now offered as part of Expert Care. Options range from a 1-year base warranty to a 5-year Expert Care Advanced or Premium.

2.2.1 DS8980F Analytic Class configuration

The DS8980F model 5341-998 is equipped with two CPCs, each with dual 11-core processors, with a maximum of 16 FC or Fibre Channel connection (IBM FICON) host adapters in the base frame and a maximum of 16 FC or FICON host adapters in a model E96 expansion frame.

Figure 2-2 shows the maximum configuration of a DS8980F with a model E96 expansion frame. The DS8950 model 996 and E96 are similar.

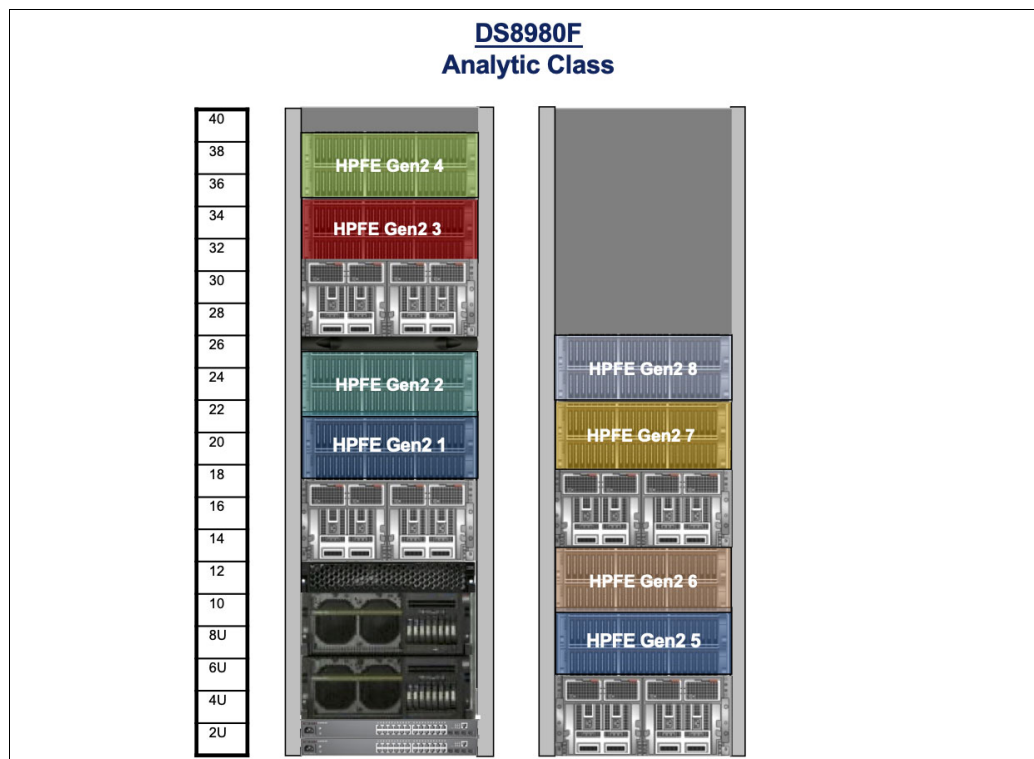


Figure 2-2 View of fully configured DS8980F or DS8950F system

Table 2-3 lists the hardware along with the minimum and maximum configuration options for the DS8980F model 998.

Table 2-3 DS8980F model 998 component list

Features	DS8980F
Rack size	19" 600x1050, 40U fixed
Frames min / max	1 / 2
CPC	Two IBM Power machine type and model (MTM) 9009-42A systems
I/O enclosure pairs min / max	1 / 4
POWER9 cores per storage system	44
System memory	4352 GB
System non-volatile storage (NVS) memory	128 GB
Host adapters min / max ^a	2 / 32
Host adapter ports min / max	8 / 128
zHyperLink adapters min / max ^b	0 / 12
Flash drives min / max	16 / 384
High-Performance Flash Enclosure (HPFE) Gen2 pairs min / max	1 / 8
Flash redundant array of independent disks (RAID) adapter pairs min / max	1 / 8
iPDU min / max	2 / 6
Power	Single-phase or three-phase
HMC	Two in ME
Ethernet switches	Two 8-port switches in ME, two extra 24-port switches with an E96 expansion frame
Keyboard/monitor	1

a. For more information, see 2.4.2, "I/O enclosure adapters" on page 51.

b. For more information, see *Getting Started with IBM zHyperLink for z/OS*, REDP-5493.

Note: The DS8900F hardware uses iPDU, non-volatile dual inline memory modules (NVDIMMs) and Backup Power Modules (BPMs) to replace the internal DC-UPSs in prior generations.

2.2.2 DS8950F Agility Class configuration

The DS8950F model 5341-996 is equipped with two CPCs, each with either single or dual 10-core processors, with a maximum of 16 FC / FICON host adapters in the base frame and a maximum of 16 FC / FICON host adapters in the expansion frame.

Figure 2-2 on page 28 shows the maximum configuration of a DS8950F model 996 and DS8950F model E96.

Table 2-4 lists the hardware along with the minimum and maximum configuration options for the DS8950F model 996.

Table 2-4 DS8950F model 996 component list

Features	DS8950F
Rack size	19" 600x1050, 40U fixed
Frames min / max	1 / 2
CPC	Two IBM Power MTM 9009-42A systems
I/O Enclosure pairs min / max	1 / 4
POWER9 cores per system min / max	20 / 40
System memory min / max	512 GB / 3456 GB
System NVS memory min / max	32 GB / 128 GB
Host adapters min / max ^a	2 / 32
Host adapter ports min / max	8 / 128
zHyperLink adapters min / max ^b	0 / 12
Flash drives min / max	16 / 384
HPFE Gen2 pairs min / max	1 / 8
Flash RAID adapter pairs min / max	1 / 8
iPDU min / max	2 / 6
Power	Single-phase or three-phase
HMC	Two in ME
Ethernet switches	Two 8-port switches in ME, two extra 24-port switches with an E96 expansion frame
Keyboard/Monitor	1

a. For more information, see 2.4.2, "I/O enclosure adapters" on page 51.

b. For more information, see *Getting Started with IBM zHyperLink for z/OS*, REDP-5493.

2.2.3 DS8910F Flexibility Class Racked configuration

The DS8910F model 5341-994 is equipped with two CPCs, each with dual quad-core processors, with a maximum of 16 Fibre Channel (FC) / FICON host adapters, and a maximum of four zHyperLink adapters.

Figure 2-3 shows a DS8910F 994 system.

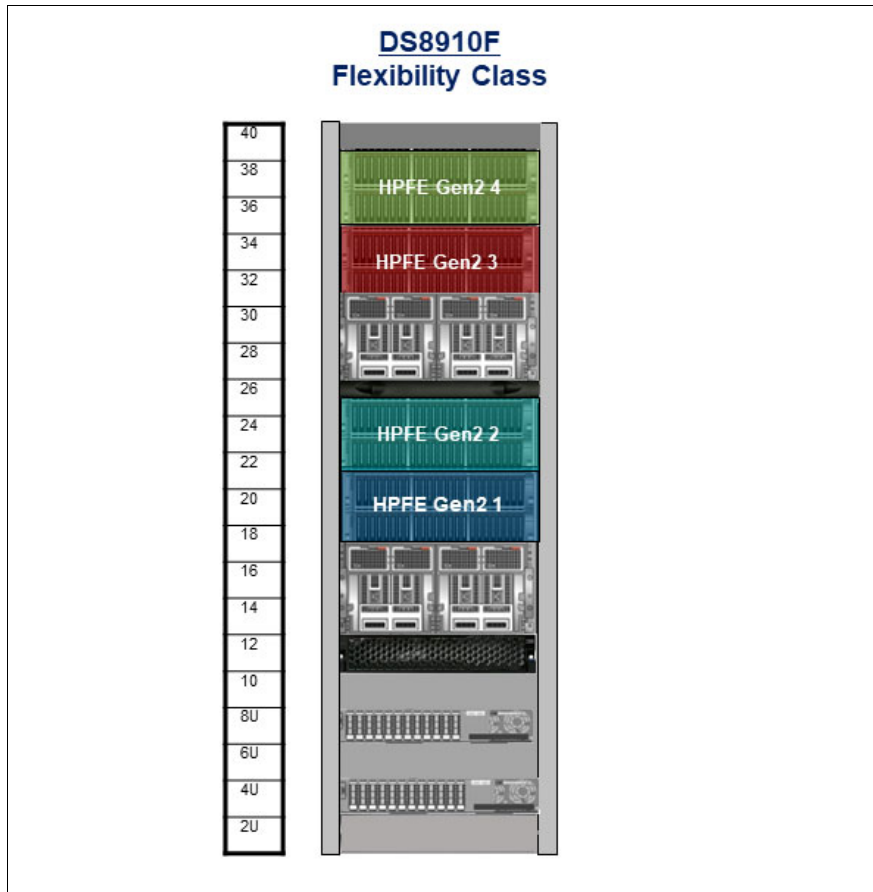


Figure 2-3 View of fully configured DS8910F model 994

Table 2-5 lists the hardware and the minimum and maximum configuration options for the DS8910F model 994.

Table 2-5 DS8910F model 994 component list

Features	DS8910F
Rack size	19" 600x1050, 40U fixed
Frames	1
CPC	Two IBM Power MTM 9009-22A systems
I/O bay pairs min / max	1 / 2
POWER9 cores per storage system	16
System memory min / max	192 GB / 512 GB
System NVS memory min / max	8 GB / 32 GB

Features	DS8910F
Host adapters min / max ^a	2 / 16
Host adapter ports min / max	8 / 64
zHyperLink adapters min / max ^b	0 / 4
Flash drives min / max	16 / 192
HPFEs Gen2 pairs min / max	1 / 4
Flash RAID adapter pairs min / max	1 / 4
iPDU min / max	2 / 4
Power	Single-phase or three-phase
HMC	2 in ME
Ethernet switches	Two 8-port switches in ME
Keyboard/Monitor	1

a. For more information, see 2.4.2, “I/O enclosure adapters” on page 51.

b. For more information, see *Getting Started with IBM zHyperLink for z/OS*, REDP-5493.

2.2.4 DS8910F Flexibility Class Rack-Mounted configuration

The DS8910F Flexibility Class Rack-Mounted model 993 provides a modular rack-mountable enterprise storage system within the 5341 all-flash family.

Note: The DS8910F Flexibility Class Rack-Mounted system has hardware specifications that differ slightly from the rack-based models. Specific information about the model 993 can be found in *IBM DS8910F Model 993 Rack-Mounted Storage System Release 9.1*, REDP-5566.

DS8910F model 993 can be integrated into existing IBM Z models T02 or ZR1 (#0937), IBM LinuxONE III model LT2 or LinuxONE II model LR1 (#0938), or any other standard 19-inch rack that conforms to EIA 310D specifications (#0939).

The model 993 uses the same hardware components that are found in the rack-based DS8910 systems and offers all the same advanced features while reducing data center footprint and power infrastructure requirements.

The DS8910F model 993 is equipped with two CPCs, each with dual quad-core processors, and it can be scaled up to 96 Tier 0, Tier 1, or Tier 2 flash drives; up to 512 GB system memory and 32 host adapter ports; and four zHyperLink adapters.

Figure 2-4 shows a DS8910F Rack-Mounted model 993 system that is integrated into the z15 model T02 or LinuxONE III model LT2 that is powered by the second set of IBM Z iPDUs. DS8910F model 993 does not share a keyboard and display with the T02 or LT2 through the IBM Z KVM. Feature Codes 0611 and 0621 are required for integration of the DS8910F model 993 into a T02 or LT2.




Figure 2-4 DS8910F model 993 integrated into a T02 or LT2

Figure 2-5 shows two DS8910F Rack-Mounted model 993 servers. One is integrated into the IBM z14® model ZR1 or LinuxONE II LR1 and is powered by the second set of IBM Z iPDUs (A3 and A4).

The console is shared with ZR1 or LR1 through the IBM Z KVM. Feature Codes 0610 and 0620 are required for integration of DS8910F model 993 into a ZR1 or LR1.

The other DS8910F Rack-Mounted model 993 server that is shown is the maximum configuration (two HPFE pairs) with a pair of optional iPDUs that can be integrated into a standard 19-inch rack. The DS8910F model 993 iPDUs support single or three-phase power.

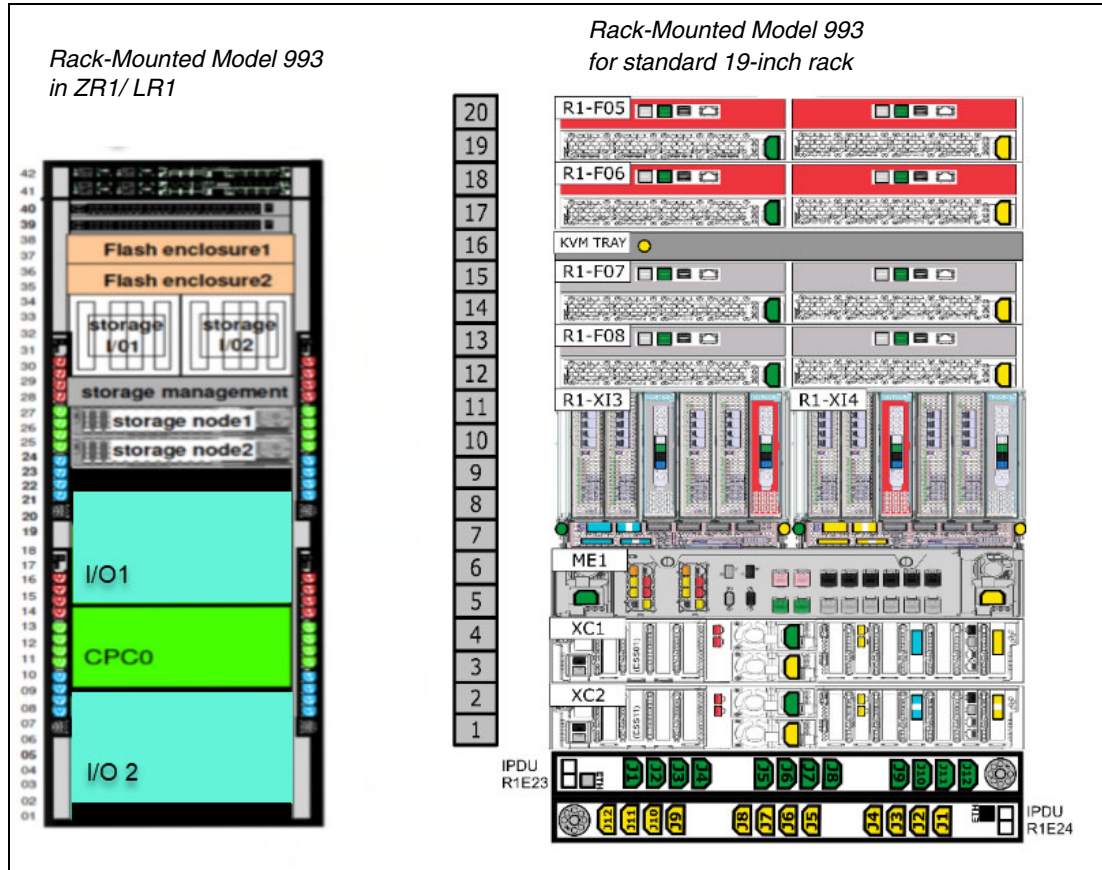


Figure 2-5 DS8910F model 993 for installation in a ZR1 or LR1 (left) and standard rack (right)

Table 2-6 lists all the hardware components and maximum capacities that are supported for the DS8910F model 993. When integrated into the entry models of IBM Z family of servers, T02, LT2, ZR1, or LR1, the DS8910F model 993 uses two IBM Z iPDUs, A3 and A4, for power. It shares the display and keyboard with ZR1 or LR1 through IBM Z KVM. It has a dedicated display and keyboard when it is integrated into a T02 or LT2.

Table 2-6 DS8910F model 993 components

Features	DS8910F model 993
Integration	T02, LT2, ZR1, LR1, or standard 19-inch rack space that conforms to EIA 310D
CPC	Two IBM Power MTM 9009-22A
I/O Enclosure pairs min / max	1 / 1

Features	DS8910F model 993
POWER9 cores per storage system min / max	16 / 16
System memory min / max	192 GB / 512 GB
System NVS memory min / max	8 GB / 32 GB
Host adapters min / max ^a	2 / 8
Host adapter ports min / max	8 / 32
zHyperLink adapters min / max	0 / 4
Flash drives min / max	16 / 96
HPFE Gen2 pairs min / max	1 / 2
Flash RAID adapter pairs min / max	1 / 2
iPDU	Two optional iPDUs with standard rack integration
Power	Single-phase or three-phase
HMC	Two in ME
Ethernet switches	Two 8-port switches in ME
Keyboard/Monitor	One with integration into T02 or LT2 and One optional with standard rack integration

a. For more information, see 2.4.2, “I/O enclosure adapters” on page 51.

2.2.5 DS8900F base frames

As mentioned in 2.1.1, “Storage system” on page 26, the DS8900F is available in different racked and rack-mounted models. The frame model number is determined by the configuration and the hardware version. The specific combinations are shown in Table 2-1 on page 26.

Note: The DS8900F models 998, 996, and 994 use a high-end 40U rack with a reduced footprint.

The DS8900F base racks accommodate the following components:

- Up to four HPFE Gen2 enclosure pairs in the base frame

The flash drives are installed in groups of 16, which are called *installation groups*. The installation groups are ordered as a set by Feature Code. Each HPFE Gen2 enclosure pair can accommodate 16, 32, or 48 flash drives.

Flash drives are available in 800 GB, 1.6 TB, 1.92 TB, 3.2 TB, 3.84 TB, 7.68 TB, and 15.36 TB capacities. For more information about HPFE, see *IBM DS8000 High-Performance Flash Enclosure Gen2 (DS8000 R9.0)*, REDP-5422.

All enclosures have redundant power and integrated cooling, which draws air from front to rear. For more information about cooling requirements, see Chapter 5, “IBM DS8900F physical planning and installation” on page 141.

► The rack power subsystem

The DS8900F family introduced a simplified rack power distribution system by using iPDUs that support single or three-phase power in all models. Each iPDU has its own dedicated input AC power cord. Various connector options are available for different regions.

The iPDUs are organized and installed in pairs. One iPDU pair is installed by default in each frame. A second iPDU pair is installed in the base frame when a second I/O enclosure or HPFE Gen2 pair are added. The second iPDU pair does not have to be the same type as the first pair. For more information about the iPDUs, see 2.6, “Power and cooling” on page 58.

► Management Enclosure

Each base frame includes a ME, which contains two HMCs in all DS8900F models. The ME also contains other essential management components, such as redundant network switches (used only for internal communications), Rack Power Control (RPC) cards, and the local or remote switch card. For more information about the HMC, see 2.7, “Management Console and network” on page 67.

Each base frame includes a 1U keyboard or display tray that is used to control both HMCs.

► POWER9 processor-based CPCs

Each base frame accommodates two POWER9 servers, also known as CPCs. The POWER9 servers contain the processors and memory that drive all functions to operate the DS8900F storage facility image (SFI). System memory and processor cores in the DS8900F can be upgraded concurrently. The Analytic and Agility Class DS8950F configurations run on two 4U IBM Power S924 servers. The Flexibility Class DS8910F systems use two 2U IBM Power S922 servers. For more information about the CPCs, see 2.3.1, “IBM POWER9 processor-based CPCs” on page 42.

Each CPC can accommodate an optional 10 Gbps Transparent Cloud Tiering (TCT) network adapter. For more information about slots for TCT, see 2.3.4, “Ethernet connections” on page 46.

► I/O enclosures

The DS8900F base frame accommodates a maximum of four I/O enclosures, which are installed in pairs. The first I/O enclosure pair is installed by default. The I/O enclosures provide PCIe Gen3 connectivity between the CPCs and installed I/O adapters.

The I/O enclosures house the following PCIe I/O adapters:

- Up to 16 host adapters in the base frame for a total of up to 64 host ports.
- Up to four 4-port 16 Gbps Fibre Channel Protocol (FCP) / FICON host adapters in each I/O enclosure, supporting either shortwave (SW) or longwave (LW).
- Up to four 4-port 32 Gbps encryption-capable FCP / FICON host adapters in each I/O enclosure, supporting either SW or LW.

Note: An intermix of 16 Gbps and 32 Gbps host adapters is supported (any combination). Intermix of SW and LW adapters is also allowed. For HA, IBM recommends installing host adapters in pairs.

The host adapter ports can be configured in the following manner:

- Switched Fibre Channel Protocol (FCP), which is used for open systems host attachment, and for Metro Mirror (MM) and Global Copy (GC).
- FICON for IBM Z host connectivity and also for z/OS Global Mirror (zGM).
- FCP and FICON are not supported simultaneously on the same port.

Note: The Fibre Channel Arbitrated Loop (FC-AL) topology is no longer supported on DS8900F host adapters.

- Flash RAID device adapters (DAs):
 - One flash RAID adapter pair is required for each HPFE Gen2 pair.
 - Each I/O enclosure pair supports up to two DA pairs.
 - DA pairs connect to HPFE Gen2 pairs over eight SAS paths in a redundant dual loop topology.
- zHyperLink connections to IBM Z hosts:
 - Supports direct connectivity to IBM Z at distances up to 150 m.
 - zHyperLink adapters (CXP transceivers) and cables are installed in pairs within an I/O enclosure pair boundary.
 - zHyperLink cables are available in lengths of 3 m (for integrated 993 models), 40 m, or 150 m. For other lengths, see *IBM Z Connectivity Handbook*, SG24-5444, or consult your optical cable vendor.
 - High-Performance FICON (z/HPF) connectivity is required.
- ▶ zHyperLink adapter and zHyperLink cables

zHyperLink connections with IBM Z hosts provide low latency for random reads and small block sequential writes. It is a point-to-point connection, as shown in Figure 2-6.

Note: All DS8900F models have zHyperLink capability.

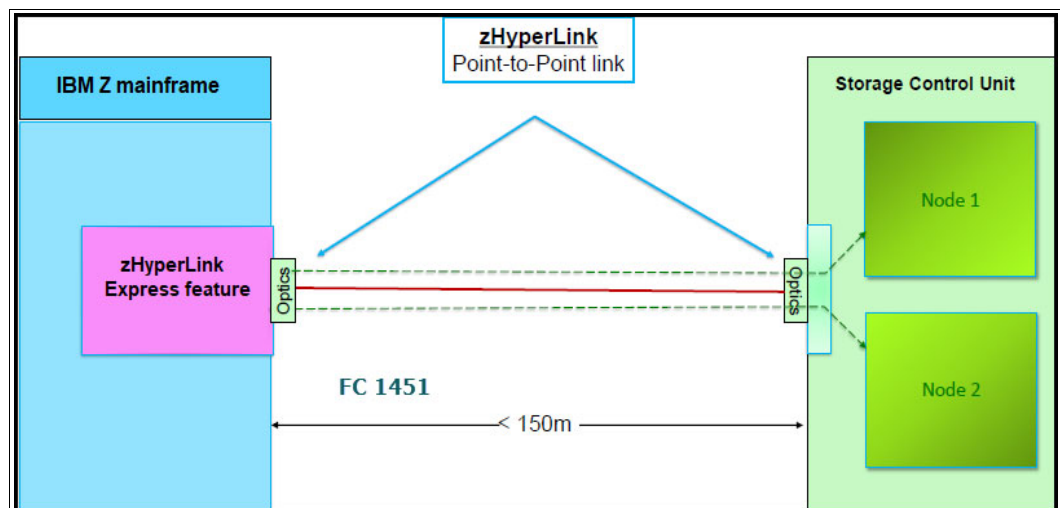


Figure 2-6 DS8900F zHyperLink connection to the system

- Each zHyperLink requires a transceiver to connect the optical zHyperLink cable to the storage system. The transceiver connects directly to an I/O enclosure CXP PCIe port, and provides connectivity for 12 transmit and receive pairs, in a multi-fiber termination push-on (MTP-24) connection. A 24x MTP-MTP cable is required for each zHyperLink connection. The transceivers are plugged into ports T3 and T4 of the I/O enclosure, as shown in Figure 2-20 on page 52.
- The number of zHyperLink adapters that can be installed into the DS8900F system depends on the number of cores per CPC. The supported combinations for zHyperLink port availability are shown in Table 2-7.

Table 2-7 zHyperLink availability for DS8900F models

System or model	Cores per CPC (DS8900F server)	zHyperLink support	Max zHyperLink connections (increments of 2)
DS8980F Base Frame model 998	22	Yes	8
DS8980F model E96		Yes	8
DS8950F Base Frame model 996	10	Yes	6
	20	Yes	8
DS8950F model E96		Yes	8
DS8910F model 994	8	Yes	4
DS8910F model 993	8	Yes	4

Note: A maximum of 12 zHyperLink connections may be configured per DS8900F system.

For more information about I/O enclosures and I/O adapters, see 2.4, “I/O enclosures and adapters” on page 48.

2.2.6 DS8900F expansion frame

The Analytic and Agility class configurations support an optional model E96 40U high-end expansion frame. The supported expansion frame options are shown in Table 2-1 on page 26.

Note: Only the DS8980F and DS8950F models support an expansion frame.

All DS8980F systems support the installation of an expansion frame without any additional features. DS8950F systems require at least 20 cores per CPC and 1 TB system memory (system cache) to support the extra throughput that is provided by the installation of I/O and storage enclosures in the expansion frame.

Expansion frame components

The model E96 expansion frame accommodates the following components:

- ▶ Up to four HPFE Gen2 enclosure pairs in the expansion frame.

The flash drives are installed in groups of 16, which are called *installation groups*. The installation groups are ordered as a set by Feature Code. Each HPFE Gen2 enclosure pair can accommodate 16, 32, or 48 flash drives.

Flash drives are available in 800 GB, 1.6 TB, 1.92 TB, 3.2 TB, 3.84 TB, 7.68 TB, and 15.36 TB capacities. For more information about HPFE, see *IBM DS8000 High-Performance Flash Enclosure Gen2 (DS8000 R9.0)*, REDP-5422.

All enclosures have redundant power and integrated cooling, which draws air from front to rear. For more information about cooling requirements, see Chapter 5, “IBM DS8900F physical planning and installation” on page 141.

- ▶ The rack power subsystem.

The DS8900F family introduced a simplified rack power distribution system that uses iPDUs that support single or three-phase power in all models. Each iPDU has its own dedicated input AC power cord. Various connector options are available for different regions.

The iPDUs are ordered and installed in pairs. One iPDU pair is installed by default in each frame. For more information about iPDUs, see 2.6, “Power and cooling” on page 58.

- ▶ I/O enclosures.

The DS8950F model E96 accommodates a maximum of four I/O enclosures, which are installed in pairs. The first I/O enclosure pair is installed by default. The I/O enclosures provide PCIe Gen3 connectivity between the CPCs and installed I/O adapters.

The I/O enclosures house the following PCIe I/O adapters:

- Up to 16 host adapters in the base frame for a total of up to 64 host ports.
- Up to four 4-port 16 Gbps FCP / FICON host adapters in each I/O enclosure, supporting either SW or LW.
- Up to four 4-port 32 Gbps FCP / FICON encryption-capable host adapters in each I/O enclosure, supporting either SW or LW.

Note: An intermix of 16 Gbps and 32 Gbps host adapters is supported (any combination). Intermix of SW and LW adapters is also allowed. For HA, IBM recommends installing host adapters in pairs.

The host adapters can be configured in the following manner:

- Switched Fibre Channel Protocol (FCP), which is used for open systems host attachment, and for Metro Mirror (MM) and GC.
- FICON for IBM Z host connectivity and also for zGM.
- FCP and FICON are not supported simultaneously on the same port.

Note: The FC-AL topology is no longer supported on DS8900F host adapters.

- zHyperLink connections to IBM Z hosts:
 - Supports direct connectivity to IBM Z at distances up to 150 m.
 - zHyperLink adapters (CXP transceivers) and cables are installed in pairs, within an I/O enclosure pair boundary.
 - zHyperLink cables are available in lengths of 3 m (for integrated 993 models), 40 m, or 150 m. For other lengths, see *IBM Z Connectivity Handbook*, SG24-5444, or contact your optical cable vendor.
- Flash RAID DAs:
 - One flash RAID adapter pair is required for each HPFE Gen2 pair.
 - Each I/O enclosure pair supports up to two DA pairs.
 - DA pairs connect to HPFE Gen2 pairs over eight SAS paths in a redundant dual loop topology.

For more information, see 2.4, “I/O enclosures and adapters” on page 48.

Connecting the expansion frame

The default cable set for connecting an expansion frame contains PCIe copper cables that allow a distance of up to 2 m. This limitation requires the installation of the expansion frame next to the base frame. One cable set is required for each installed I/O enclosure pair in the expansion frame.

To ease floor planning for future expansions, an available optical PCIe cable allows a distance up to 20 m. The cable set contains optical cables and transceivers. One cable set is required for each installed I/O enclosure pair in the expansion frame.

As shown in Figure 2-7 on page 41, this extension makes the positioning of an expansion frame more flexible, especially for future capacity expansion. An extra rack side cover pair is available if needed.

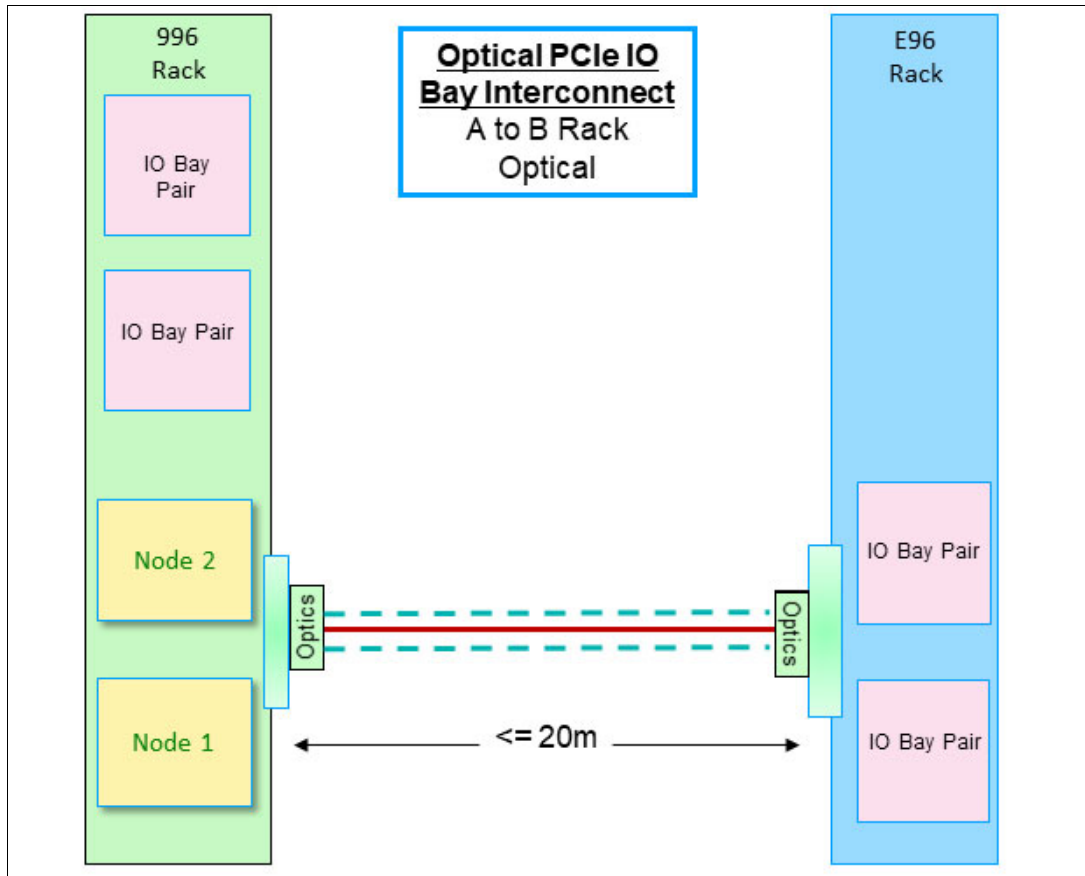


Figure 2-7 Expansion frame at a distance

2.2.7 Scalable upgrades

The hardware features that are supported in the DS8900F depend on the total system memory and total processor cores that are installed. This design ensures that the performance and capacity scale correctly. For more information about processor and system memory requirements for hardware upgrades, see Figure 2-11 on page 44. Each of the DS8900F configurations can be nondisruptively upgraded from the smallest system memory feature to the largest memory feature that is supported by that configuration.

2.2.8 Licensed functions

Several of the DS8900F functions require a license key. For more information about licensed functions, see Chapter 7, “IBM DS8900F features and licensed functions” on page 199.

2.3 DS8900F architecture overview

This section provides an architectural overview of the major components of the DS8900F:

- ▶ IBM POWER9 processor-based CPCs
- ▶ I/O enclosures and adapters
- ▶ PCIe connectivity and communication
- ▶ Storage subsystem
- ▶ Hardware management

2.3.1 IBM POWER9 processor-based CPCs


All DS8900F systems include two CPCs:

- ▶ In the DS8980F and DS8950F configurations, the CPCs are IBM Power 9009-42A servers, which have two processor sockets. The processors run in Turbo mode and typically achieve cycle speeds of 3.5 GHz.
 - For the DS8980F configuration, the CPCs are populated with two 11-core processors for a combined total of 22 cores per CPC. The DS8980F CPCs support a total of 2176 GB of memory (two modules of 32 GB DDR4 NVDIMMs, two modules of 32 GB DDR4 RDIMMs, and 16 modules of 128G DDR4 RDIMMs).
 - For the DS8950F configuration, the CPCs are populated with one 10-core processor in the SCM-0 slot, or with two 10-core processors, for a combined total of 20 cores per CPC. The DS8950F CPCs support a maximum of 1728 GB of memory (two modules of 32 GB DDR4 NVDIMMs and 26 modules of 64 GB DDR4 RDIMMs).
- ▶ The CPCs that are used in DS8980F and DS8950F feature the following configuration:
 - Two single-chip module (SCM) sockets.
 - Thirty-two DDR4 dual inline memory module (DIMM) slots.
 - Two BPMs.
 - Three PCIe Gen4 16 lane slots with 16 lanes.
 - Two PCIe Gen4 16 lane slots with eight lanes.
 - Two PCIe Gen3 16 lane slots with eight lanes.
 - Four PCIe Gen3 eight lane slots with eight lanes.
 - One storage cage with two hard disk drives (HDDs).
 - Four power supply units (PSUs).

Figure 2-8 on page 43 and Figure 2-9 on page 44 show the CPC configured for the DS8980F and DS8950F systems.



Figure 2-8 DS8980F CPC: front view



DS8980F CPC: rear view

- ▶ In the DS8910F configuration, the CPCs are IBM Power 9009-22A servers, which have two processor sockets. They are populated with two 4-core processors for a combined total of eight cores per CPC. The processors run in Turbo mode and typically achieve cycle speeds of 3.4 GHz.

The DS8910F CPCs support a maximum of 256 GB of memory (two modules of 16 GB DDR4 NVDIMMs and 14 modules of 16 GB DDR4 RDIMMs).

- ▶ The CPCs that are used in DS8910F systems feature the following configuration:
 - Two SCM sockets.
 - Thirty-two DDR4 DIMM slots.
 - One BPM.
 - Three PCIe Gen4 16 lane slots with 16 lanes.
 - Two PCIe Gen4 16 lane slots with eight lanes.
 - Two PCIe Gen3 16 lane slots with eight lanes.
 - Two PCIe Gen3 eight lane slots with eight lanes.
 - One storage cage with two HDDs.
 - Two PSUs.

Figure 2-9 and Figure 2-10 show the CPC as configured in the DS8910F system.



Figure 2-9 DS8910F CPC: front view




Figure 2-10 DS8910F CPC: rear view

For more information about the server hardware that is used in the DS8910F and DS8950F, see *IBM Power Systems S922, S914, and S924 Technical Overview and Introduction*, REDP-5497.

In the DS8900F, processor core and system memory configurations dictate the hardware that can be installed in the storage system. Processors and memory can be upgraded concurrently as required to support storage system hardware upgrades. The supported maximum system hardware components depend on the total processor and system memory configuration.

Figure 2-11 shows the supported components for the DS8900F processor and memory options. NVS values are typically 1/16th of installed system memory, except for the smallest systems with 192 GB system memory, where only 8 GB, that is, 4 GB per CPC, is used as NVS, and for the biggest systems of 3.4 or 4.3 TB memory, where NVS remains at 128 GB.

Model	Processor cores per CPC	Processor sockets per CPC	SMT config	Max available threads	Max zHyperLink threads	System memory (GB)	NVS (GB)	Expansion Frame	Max HA (ports)	Max I/O Enclosure pairs	Max zHyperLinks	Max HPFE Gen 2 pairs	Max Flash Drives
DS8980 - Analytic Class													
998	22-core	2x 11-core	4	88	22	4352	128	0-1	32 (128 ports)	4	12	8	384
DS8950F - Agility Class													
996	10-core	1x 10-core	4	40	10	512	32	0	16 (64 ports)	2	6	4	192
996	20-core	2x 10-core	4	80	20	1024 2048 3456	64 128 128	0-1	32 (128 ports)	4	12	8	384
DS8910F - Flexibility Class													
994	8-core	2x 4-core	4	32	8	192 512	8 32	0	8 (32 ports)	2	4	4	192
DS8910F - Flexibility Class Rackless													
993	8-core	2x 4-core	4	32	8	192 512	8 32	0	8 (32 ports)	1	4	2	96

Figure 2-11 Supported components for the DS8900F processor and memory options

2.3.2 Processor memory

The DS8980F configuration comes standard with 4352 GB of total system memory. The DS8950F configuration offers up to 3456 GB of total system memory. The DS8910F Racked and Rack-Mounted configurations offer up to 512 GB of total system memory.

Each CPC contains half of the total system memory. All memory that is installed in each CPC is accessible to all processors in that CPC. The absolute addresses that are assigned to the memory are common across all processors in the CPC. The set of processors is referred to as a symmetric multiprocessor (SMP) system.

The POWER9 processor that is used in the DS8900F operates in simultaneous multithreading (SMT) mode, which runs multiple instruction streams in parallel. The number of simultaneous instruction streams varies according to processor and LIC level. SMT mode enables the POWER9 processor to maximize the throughput of the processor cores by processing multiple concurrent threads on each processor core.

The DS8900F configuration options are based on the total installed memory, which in turn depends on the number of installed and active processor cores.

The DS8980F configuration comes standard with 22 cores per server, and 4.3 TB of total system memory. No processor core or system memory upgrades are supported at the time of writing.

The following DS8950F configuration upgrades can be performed nondisruptively:

- ▶ Processor configuration upgrade from 10 cores per server to 20 cores per server
- ▶ Memory upgrade from 256 GB per server to 512 GB per server or 1024 GB per server
- ▶ Memory upgrade from 512 GB per server to 1024 GB per server or 1728 GB per server

The following DS8910F configuration upgrade can be performed nondisruptively: Memory upgrade from 96 GB per server to 256 GB per server.

Note: System memory and processor upgrades are tightly coupled. They cannot be ordered or installed independently from each other.

Caching is a fundamental technique for reducing I/O latency. Like other modern caches, the DS8900F system contains volatile memory (RDIMM) that is used as a read/write cache, and NVDIMM that is used for a persistent memory write cache. (A portion of the NVDIMM capacity is also used for read/write cache.) The NVDIMM technology eliminates the need for the large backup battery sets that were used in previous generations of DS8000. If power is lost, the system shuts down in 20 ms, but power is maintained to the NVDIMMs, and data in the NVS partition is hardened to onboard NAND flash.

NVS scales according to the processor memory that is installed, which also helps to optimize performance. NVS is typically 1/16th of installed CPC memory, with a minimum of 8 GB and a maximum of 128 GB.

2.3.3 Flexible service processor

Each POWER9 processor-based CPC is managed by a flexible service processor (FSP). The FSP is an embedded controller that is based on an IBM PowerPC® processor.

The FSP controls power and cooling for the CPC. The FSP performs predictive failure analysis (PFA) for installed processor hardware, and performs recovery actions for processor or memory errors. The FSP monitors the operation of the firmware during the boot process, and can monitor the OS for loss of control and take corrective actions.

2.3.4 Ethernet connections

Each POWER9 processor-based CPC has a single 4-port, 1 Gbps Ethernet adapter installed. On all models, the top two of these ports are connected to the internal network switches that are described in Figure 2-31 on page 68. The bottom two RJ45 copper ports can be used for TCT.

A pair of optional adapters is available for TCT as a chargeable Feature Code. Each adapter provides two 10 Gbps small form-factor pluggable plus (SFP+) optical ports for short distances, and a pair of 1 Gbps (RJ45 copper) connectors. The standard 1 Gbps Ethernet adapter is in the same slot (P1-C11) in DS8980F, DS8950F, and DS8910F systems. The optional 10 Gbps Ethernet adapter is in the P1-C10 slot for DS8980F and DS8950F systems, and is in the P1-C4 slot in DS8910F systems. For more information about TCT, see *IBM DS8000 Transparent Cloud Tiering (DS8000 Release 9.2)*, SG24-8381.

Figure 2-12 shows the location codes of the CPCs in DS8980F and DS8950F systems. Figure 2-13 on page 47 shows the location codes of the CPC in a DS8910F system.

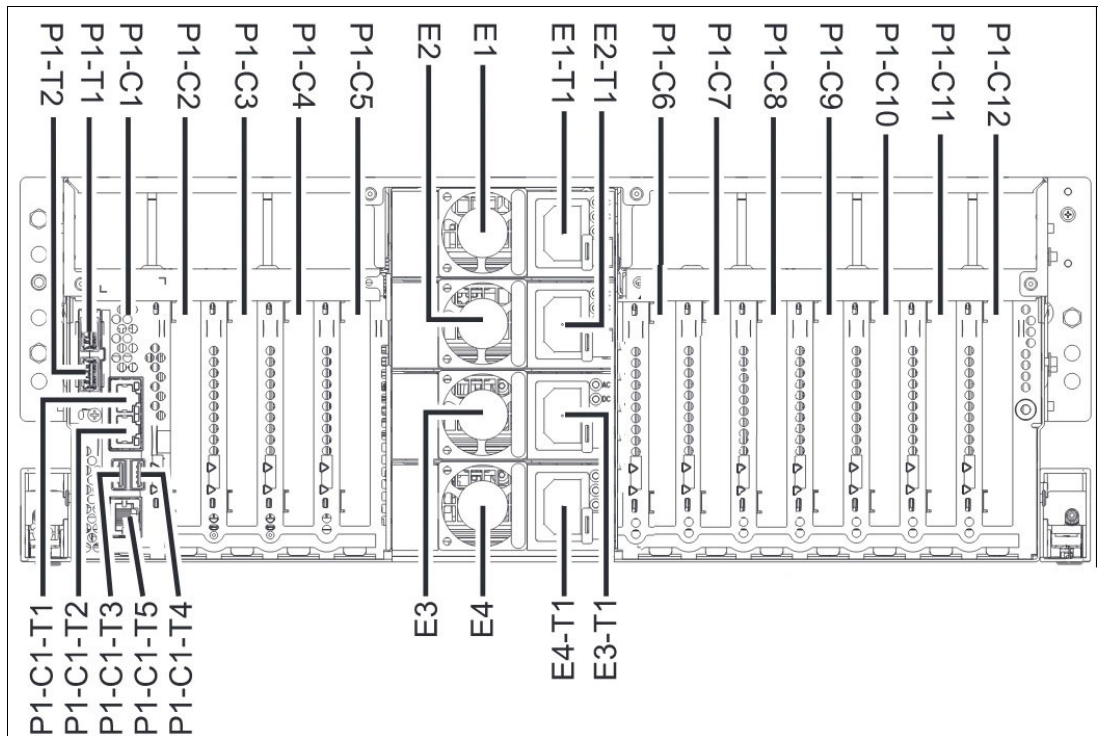


Figure 2-12 Location codes of the CPC in DS8980F and DS8950F systems in the rear




Figure 2-13 Location codes of the CPC in the DS8910F system in the rear

2.3.5 Peripheral Component Interconnect Express adapters

Each DS8900F CPC contains multiple PCIe adapters. These adapters enable point-to-point connectivity between CPCs, and I/O enclosures. Depending on the configuration, up to seven PCIe adapters are in each DS8900F CPC.

A DS8900F CPC is equipped with the following PCIe adapters:

- ▶ The DS8980F system has six single-port PCIe adapters in slots C2, C4, C6, C8, C9, and C12, and one 2-port PCIe adapter in slot C3®.
- ▶ The DS8950F system has two different processor configurations for PCIe adapters:
 - Single 10-core CPC maximum configuration for one frame has four single-port PCIe adapters in slots C6, C8, C9, and C12.
 - Dual 10-core CPC maximum configuration for two frames has six single-port PCIe adapters in slots C2, C4, C6, C8, C9, and C12, and one 2-port PCIe adapter in slot C3.
- ▶ The DS8910F model 994 has four single-port PCIe adapters in slots C6, C8, C9, and C12.
- ▶ The DS8910F model 993 has two single-port PCIe adapters in slots C6 and C12.

Figure 2-14 shows the PCIe adapter locations in the DS8980F CPC. Figure 2-15 shows the PCIe adapter locations in the DS8910F CPC.




Figure 2-14 PCIe adapter locations in the DS8980F and DS8950F CPC: rear view




Figure 2-15 PCIe adapter locations in the DS8910F CPC: rear view

2.4 I/O enclosures and adapters

The DS8900F base frame and expansion frame (if installed) contain I/O enclosures, which are installed in pairs.

The I/O enclosures are PCIe Gen3-capable, and are attached to the CPCs with 8-lane PCIe Gen3 cables. The I/O enclosures have six PCIe adapter slots, plus six CXP connectors.

- ▶ DS8980 and DS8950F CPCs have up to six 1-port and one 2-port PCIe adapters that provide connectivity to the I/O enclosures.
- ▶ DS8910F CPCs have up to four 1-port PCIe adapters that provide connectivity.

Figure 2-16 on page 49 - Figure 2-18 on page 50 show the DS8900F CPC to I/O enclosure connectivity.

The DS8980 configuration requires no extra features to support an expansion frame.

The DS8950F configuration requires two 10-core processors per CPC and 1 TB system memory to support an expansion frame.

One or two I/O enclosure pairs can be installed in the base frame of the DS8900F and also in the E96 expansion frame. Each I/O enclosure can have up to four host adapters. A maximum of 16 host adapter ports are supported in a single I/O enclosure. The I/O enclosure has two zHyperLink connections. For more information about zHyperLink availability for DS8900F models, see Table 2-7 on page 38.

Note: Model 993 supports one pair of I/O enclosures.

Each I/O enclosure has the following characteristics:

- ▶ Half-width 5U rack-mountable enclosure
- ▶ Six PCIe slots (four for host adapters, two for flash RAID adapters)
- ▶ Two PCIe connections to the CPCs
- ▶ Two zHyperLink connections
- ▶ Two PCIe Gen3 connections (unused)
- ▶ Redundant power and cooling

Figure 2-16 shows the DS8980F and DS8950F CPC to I/O enclosure connectivity.

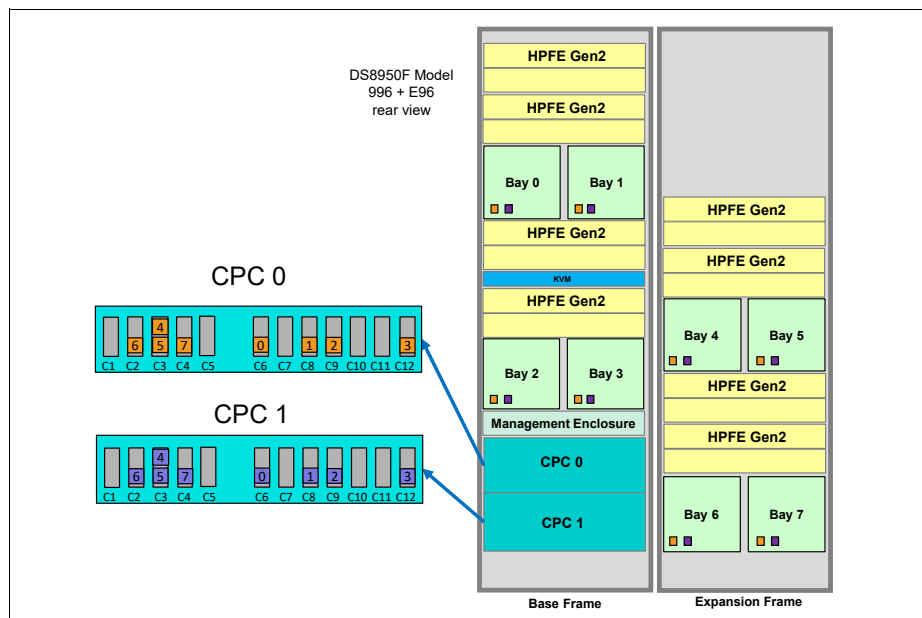


Figure 2-16 DS8980F and DS8950F I/O enclosure connections to the CPCs

Figure 2-17 shows the DS8910F model 994 CPC to I/O enclosure connectivity.

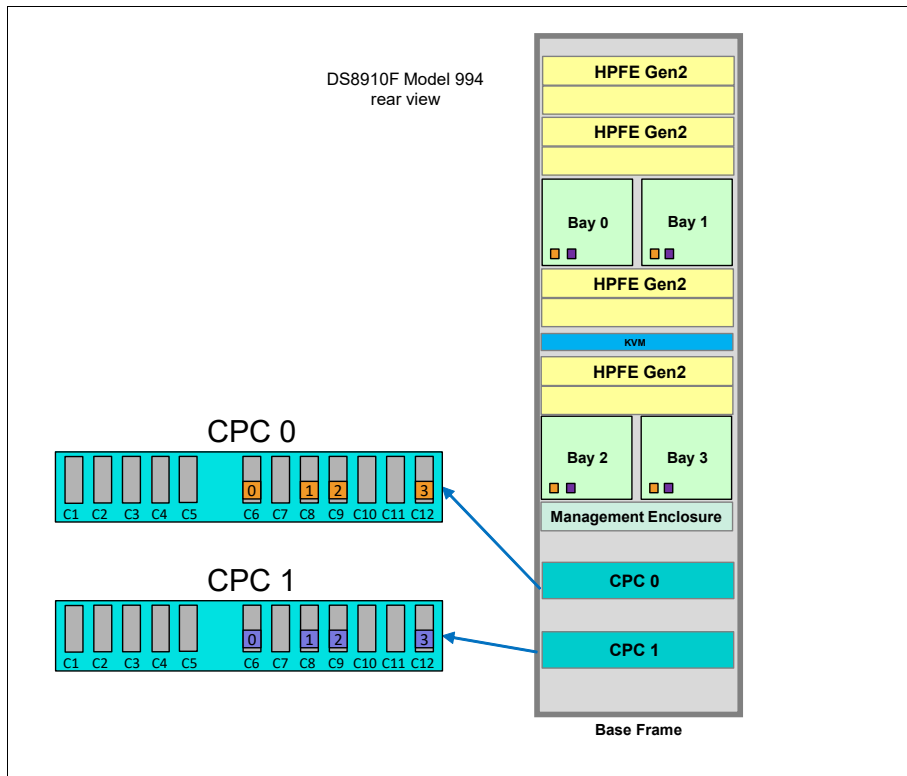


Figure 2-17 DS8910F model 994 I/O enclosure connections to the CPC

Figure 2-18 shows the DS8910F model 993 CPC to I/O enclosure connectivity.

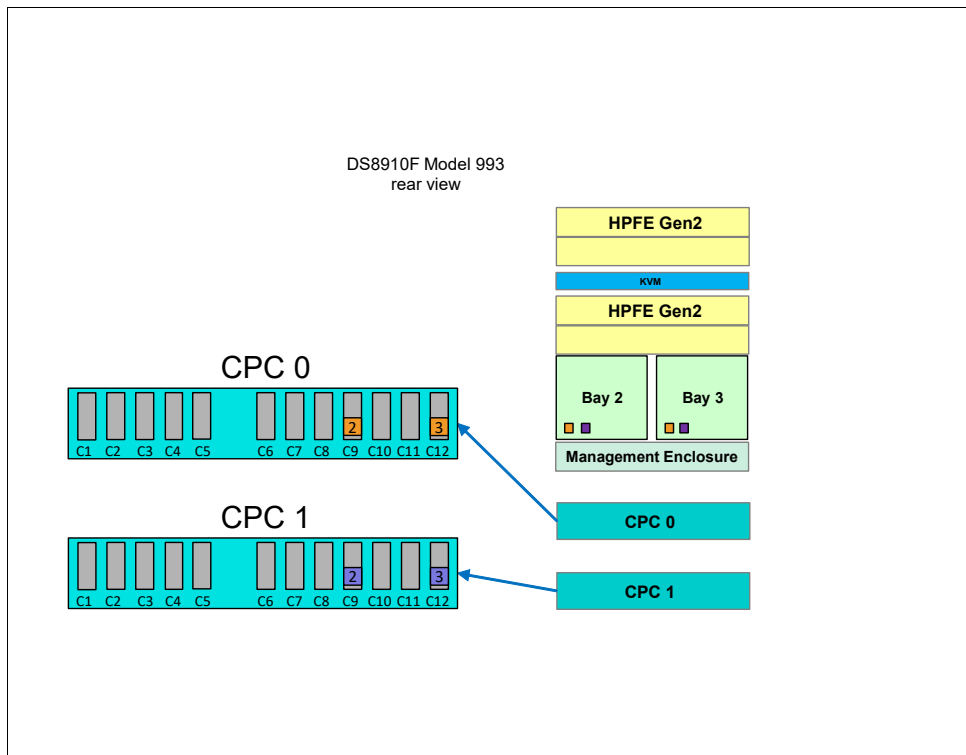


Figure 2-18 DS8910F model 993 I/O enclosure connections to the CPC

2.4.1 Cross-cluster communication

Figure 2-19 shows how the DS8900F I/O enclosure hardware is shared between the servers. One CPC is on the left side and one CPC is on the right side, and the diagram shows the SCMs (SCM #0 and SCM#1). The solid lines denote primary PCIe paths, and the dashed lines denote secondary PCIe paths.

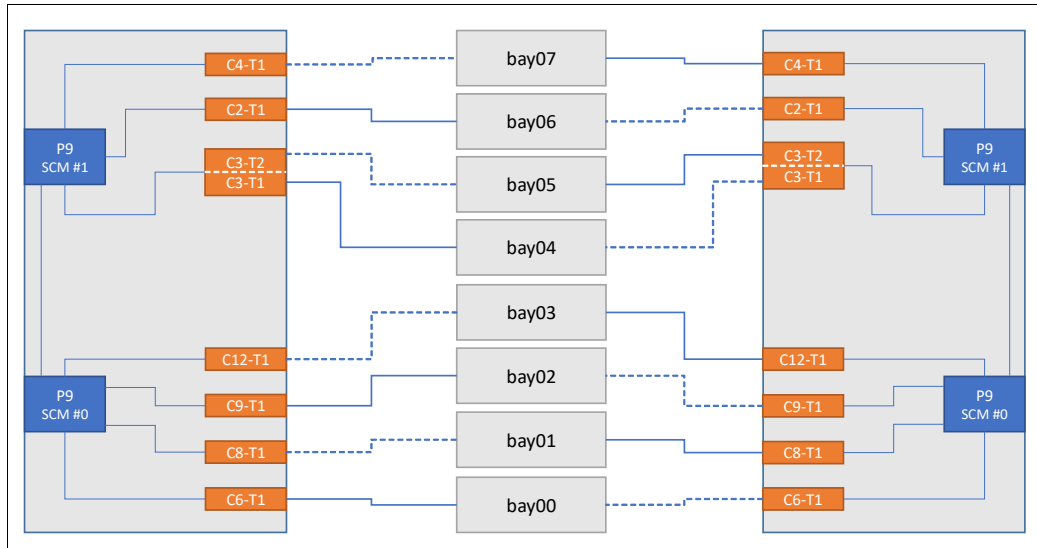


Figure 2-19 DS8900F series PCIe communications paths

The DS8900F uses the PCIe paths through the I/O enclosures to provide high-speed communication paths between the CPCs. Normally, the lowest available even-numbered I/O enclosure is used for communication from server 0 to server 1, and the lowest available odd-numbered I/O enclosure is used for communication from server 1 to server 0.

If a failure occurs in one or more I/O enclosures, any of the remaining enclosures can be used to maintain communication between the servers.

2.4.2 I/O enclosure adapters

The DS8900F I/O bay provides the connectivity from the host systems to the storage arrays through the CPCs. Each I/O adapter is optimized for its specific task.

The I/O bay can contain up to four host adapters that provide attachment to host systems and up to two flash RAID DAs to provide attachment to the HPFE Gen2 enclosures. Each I/O bay has six PCIe x8 CXP connectors on the I/O bay PCIe module. Two ports (T1 and T2) are for the internal PCIe fabric connections to CPC 0 and CPC 1. Two ports (T3 and T4) are for attachment of zHyperLink to IBM Z and two ports (T5 and T6) are unused.

Figure 2-20 shows the DS8900F I/O bay adapter layout.




Figure 2-20 DS8900F I/O bay adapter layout

DS8900F host adapters

Attached host servers interact with software that is running on the CPCs to access data that is stored on logical volumes. The CPCs manage all read/write requests to the logical volumes on the storage arrays.

Two different types of host adapters are available: 32 Gbps and 16 Gbps. Both have four ports. The 32 Gbps adapters can auto-negotiate their data transfer rate down to 8 Gbps full-duplex data transfer. The 16 Gbps adapters can auto-negotiate down to 4 Gbps full-duplex data transfer.

Figure 2-21 on page 53 shows the 32 Gbps FCP or FICON host adapter. It provides faster single stream and per-port throughput and reduces latency compared to the 16 Gbps adapter. The 32 Gbps host adapter is equipped with a quad-core 2 GHz PowerPC processor that delivers dramatic (2 - 3 times) full adapter I/O operations per second (IOPS) improvements compared to the 16 Gbps adapter. The 32 Gbps adapter is required to enable IBM Fibre Channel Endpoint Security encryption.



Figure 2-21 32 Gbps FCP or FICON host adapter

The 16 Gbps host adapter supports only IBM Fibre Channel Endpoint Security authentication.

The 32 Gbps host adapter supports both IBM Fibre Channel Endpoint Security authentication and line-rate encryption.

For more information, see *IBM Fibre Channel Endpoint Security for IBM DS8900F and IBM Z, SG24-8455*.

Both adapters contain a high-performance application-specific integrated circuit (ASIC). To ensure maximum data integrity, it supports metadata creation and checking. Each FC port supports a maximum of 509 host login IDs and 1,280 paths. This configuration enables the creation of large storage area networks (SANs).

Each host adapter port can be configured as either FICON or FCP. For both host adapters, the adapter optics can be either LW or SW.

The DS8980F and DS8950F configurations support a maximum of 16 host adapters in the base frame and 16 extra host adapters in the model E96 expansion frame. The DS8910F model 994 configuration supports a maximum of 16 host adapters. The DS8910F model 993 configuration supports a maximum of eight host adapters.

Host adapters are installed in slots 1, 2, 4, and 5 of the I/O enclosure. Figure 2-20 on page 52 shows the locations for the host adapters in the DS8900F I/O enclosure. The system supports an intermix of both adapter types up to the maximum number of ports, as shown in Table 2-8.

Optimum availability: To obtain optimum availability and performance, one host adapter must be installed in each available I/O enclosure before a second host adapter is installed in the same enclosure.

Table 2-8 DS8900F port configurations

Model	Min/Max host adapters	Min/Max host adapter ports	Max zHyperLink adapters
998	2/16	8/64	8
998+E96	2/32	8/128	12 ^a
996	2/16	8/64	6/8 ^b
996 + E96	2/32	8/128	12 ^a
994	2/16	8/64	4

a. Maximum of eight in either the base frame or the expansion frame. Maximum of 12 per system.

b. Maximum of six with 10-core processors per CPC, and eight with 20-core processors per CPC.

Table 2-9 shows the preferred host adapter installation order for the DS8900F system. The host adapter locations and installation order for the four I/O enclosures in the base frame are the same for the I/O enclosures in the first expansion frame.

Table 2-9 DS8900F host adapter installation order

I/O bay	Slot number					
	C1	C2	C3	C4	C5	C6
For two I/O enclosures (all models)						
Bottom I/O bay 02 / 06	3	7		1	5	
Bottom I/O bay 03 / 07	2	6		4	8	
For four I/O enclosures (Model 998, Model 996, Model 994 <i>Model E96^a</i>)						
Top I/O bay 00 / 04	7	15		3	11	
Bottom I/O bay 02 / 06	5	13		1	9	
Top I/O bay 01 / 05	4	12		8	16	
Bottom I/O bay 03 / 07	2	10		6	14	

a. For the DS8950F model E96, the enclosure numbers are in *emphasized* text, and the plug order is the same as the other models.

Fibre Channel

The DS8900F uses the Fibre Channel Protocol (FCP) to transmit Small Computer System Interface (SCSI) traffic inside FC frames. It also uses FC to transmit FICON traffic for IBM Z I/O.

Each of the ports on a DS8900F host adapter can be configured for FCP or FICON, but a single port *cannot* be configured for both concurrently. The port topology can be changed by using the DS GUI or DS CLI.

Fibre Channel-supported servers

The current list of servers that are supported by FC attachment can be found at the [IBM System Storage Interoperation Center \(SSIC\) website](#).

Fibre Channel distances

All ports on each adapter must be either LW or SW. The two types *cannot* be intermixed within a single adapter. With LW, you can connect nodes at distances of up to 10 km (6.2 miles) non-repeated. With SW, you are limited to a distance that depends on the FC cable type and the data transfer rate. For 16 Gbps and 32 Gbps, use OM3 or OM4. For the link distance limitations, see Table 2-10.

Table 2-10 SW link distance

Speed	OM3 link distance	OM4 link distance
16 Gbps	100 m (328 ft.)	125 m (410.1 ft.)
32 Gbps	70 m (229.6 ft.)	100 m (328 ft.)

Flash RAID adapters

Flash RAID adapters, also known as DAs, provide redundant access to the internal storage devices. Each DA manages a pair of HPFE Gen2 enclosures. The adapters are always installed as a pair. Logical configuration is then balanced across the DA pair for load-balancing and the highest throughput.

The DAs are installed in the I/O enclosures and are connected to the CPCs through the PCIe network. The DAs are responsible for managing and monitoring the flash RAID arrays. The DAs provide remarkable performance because of a high-function and high-performance ASIC. To ensure maximum data integrity, the adapter supports metadata creation and checking.

For more information about the flash RAID adapters, see *IBM DS8000 High-Performance Flash Enclosure Gen2 (DS8000 R9.0)*, REDP-5422.

2.5 Flash drive enclosures

The DS8900F is equipped with HPFE Gen2 storage enclosures:

- ▶ Flash enclosure pairs connect to flash RAID DAs by using SAS cabling.
- ▶ Flash RAID DAs are installed in the C3 and C6 slots of I/O enclosures.

HPFE Gen2 enclosures are always installed in pairs. Each enclosure pair supports 16, 32, or 48 flash drives. A single Gen2 enclosure is shown in Figure 2-22.



Figure 2-22 HPFE Gen2 enclosure

Each HPFE Gen2 pair is connected to a redundant pair of Flash-optimized RAID controllers. The PCIe flash RAID adapters are installed in the DS8900F I/O enclosures.

The DS8980F and DS8950 configurations can support up to four HPFE Gen2 pairs in the base frame, and up to four HPFE Gen2 pairs in the expansion frame for a total of eight HPFE Gen2 pairs, with a maximum of 384 flash drives.

To learn more about the HPFE Gen2, see *IBM DS8000 High-Performance Flash Enclosure Gen2 (DS8000 R9.0)*, REDP-5422.

Flash drives in the HPFE Gen2

Each HPFE Gen2 pair can contain 16, 32, or 48 flash drives. Flash drives are available in 800 GB, 1.6 TB, 1.92 TB, 3.2 TB, 3.84 TB, 7.68 TB, or 15.36 TB capacities. All flash drives in an HPFE Gen2 enclosure pair must be of the same type (high performance or high capacity).

Flash drive sets

Flash drives are ordered in drives sets of 16. The HPFE Gen2 pair can contain 16, 32, or 48 flash drives (1, 2, or 3 drive sets). Half of the drive set is installed in each enclosure of the pair.

Storage-enclosure fillers

Storage-enclosure fillers occupy empty drive slots in the storage enclosures. The fillers ensure consistent airflow through an enclosure. For HPFE Gen2, one filler feature provides a set of 16 fillers.

High-performance flash drives

The DS8900F system supports 2.5-inch high-performance flash drives, which are designated as Flash Tier 0 (see Table 2-11). All high-performance flash drives are Full Disk Encryption (FDE) capable. For more information about licensed features, see Chapter 7, “IBM DS8900F features and licensed functions” on page 199.

Table 2-11 Supported high-performance flash drives

Feature Code	Drive capacity	Drive type	RAID support (Default RAID 6)
1611	800 GB	2.5-in flash tier 0	5, 6, and 10 ^a
1612	1.6 TB	2.5-in Flash Tier 0	6 and 10
1613	3.2 TB	2.5-in Flash Tier 0	6 and 10

a. RAID 5 is supported, but not recommended.

Note: To learn more about the DS8900F drive features, see the *IBM System Storage DS8900F Introduction and Planning Guide*, SC27-9560.

High-capacity flash drives

The DS8900F system also supports 2.5-inch high-capacity flash drives (see Table 2-12). All high-capacity flash drives are FDE-capable. 3.84 TB drives are designated as Flash Tier 1, while 1.92 TB, 7.68 TB, and 15.36 TB drives are designated as Flash Tier 2.

Table 2-12 Supported high-capacity flash drives

Feature Code	Drive capacity	Drive type	RAID support (Default RAID 6)
1623	3.84 TB	2.5-in. Flash Tier 1	6, 10
1622	1.92 TB	2.5-in. Flash Tier 2	6, 10
1624	7.68 TB	2.5-in. Flash Tier 2	6, 10
1625	15.36 TB	2.5-in. Flash Tier 2	6, 10

Arrays and spares

Each HPFE Gen2 pair can contain up to six array sites. Each set of 16 flash drives creates two 8-drive array sites. During logical configuration, RAID 6 arrays are created by default on each array site, and the required number of spares are created. Each HPFE Gen2 pair always has two global spares, which are created from the first increment of 16 flash drives. For RAID 6 arrays, the first two arrays that are created from these array sites are 5+P+Q+S. Subsequent RAID 6 arrays in the same HPFE Gen2 Pair are 6+P+Q.

Note: For all drive types, RAID 6 is the default in DS GUI and DS CLI, but RAID 10 is optional. For flash drives smaller than 1 TB, RAID 5 is also optional, but is not recommended.

System capacity limitations

There are capacity limitations depending on the use of small or large extents. The maximum amount of usable and provisioned capacities from small and large extents depend on the amount of system cache, as shown in Table 2-13.

Table 2-13 Maximum usable and provisioned capacity based on system cache size

Cache	Max. usable size with large extents	Max. provisioned size with large extents	Max. usable size with small extents	Max. provisioned size with small extents
Less than or equal to 512 GB	Fixed-Block (FB): 4096 TiB Count Key Data (CKD): 3652 TiB	FB: 4096 TiB CKD: 3652 TiB	FB: 512 TiB CKD: 551 TiB	FB: 1024 TiB CKD: 913 TiB
Greater than 512 GB	FB: 16384 TiB CKD: 14608 TiB	FB: 8160 TiB - 16384 TiB ^a CKD: 7263 TiB - 14608 TiB ^a	FB: 2048 TiB CKD: 2205 TiB	FB: 3968 TiB - 4096 TiB ^a CKD: 3538 TiB - 3652 TiB ^a

a. The exact value within the range is determined by a complex calculation that is based on the number of volumes and volume sizes. You should conservatively plan for configurations targeting the low end of the range.

Table 2-14 shows the maximum number of flash drives and maximum raw storage capacity for the different models.

Table 2-14 Maximum raw storage capacity per model

Model/Processors	System memory (GB)	Maximum flash drives	Maximum raw storage capacity ^a
DS8980F / 22-core	4 352	384	5 898 TB
DS8950F / 10-core	512	192	2 949 TB
DS8950F / 20-core	1 024 2 048 3 456	384	5 898 TB
DS8910F / 8-core	192 512	192	2 949 TB

a. Using 15.36 TB Flash Tier 2 drives.

2.6 Power and cooling

The DS8900F power and cooling systems are highly redundant. The components are described in this section. For more information, see 3.6, “RAS on the power subsystem” on page 97.

2.6.1 Rack power control cards

The DS8900F features a pair of redundant rack power control cards (RPCs), which monitor hardware conditions and provide control paths to the I/O enclosures and LED indicators in the storage system. The RPCs are housed in the ME and get their power from the ME redundant PSUs. The RPCs are hardware-based system management controllers with their own firmware and code.

As in earlier DS8000 models, the DS8900F RPCs are connected to the FSPs in each CPC by using serial Inter-Integrated Circuit (I²C) cables and to all I/O enclosures in the base and expansion rack over daisy-chained Power Control Network (PCN) connections.

RPCs also communicate with each of the CPC operating LPARs over RS485 serial connections. Using this communication path, the RPCs act as a quorum in the CPC or LPAR cluster communication to avoid cluster splits in a quorum or RPC race.

Important: Unlike earlier DS8000 models, DS8900F RPCs normally provide for only communication and connectivity to components in the storage system. Power control functions are managed by the HMCs.

2.6.2 Intelligent Power Distribution Units

The usage of NVDIMMs for write cache retention allows DS8900F to greatly simplify rack power distribution. Bulky DC-UPS battery backup systems are replaced by compact intelligent power distribution units (iPDUs) that significantly reduce the rack footprint and weight.

The iPDUs are available in single or three-phase power configurations in all models. Each iPDU has one AC power connector and a dedicated inline power cord. Output power is provided by 12 C13 power outlets with circuit breaker protection.

Figure 2-23 shows the iPDU connections.

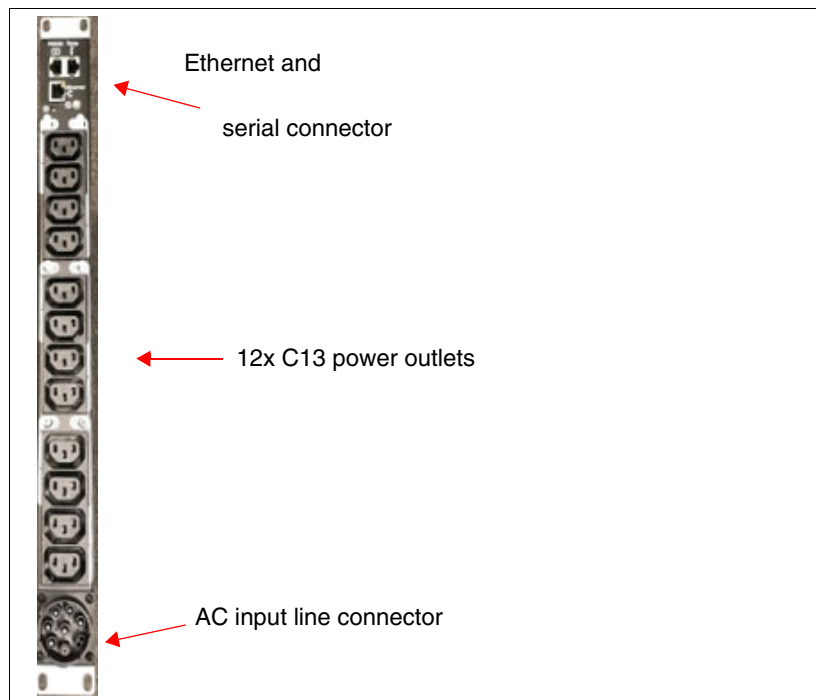


Figure 2-23 iPDU connections

iPDUs are installed in pairs. Each DS8900 rack has a minimum of one iPDU pair. For models 988, 986, and 984, a second pair may be installed in the base frame to provide power for more I/O and storage enclosures.

DS8900F can tolerate a power line disturbance (PLD) of up to 20 ms. If the PLD exceeds this threshold on both sides of the power domain, the system initiates an orderly shutdown, and data in write cache is saved to flash memory in the NVDIMMs. The NVDIMMs remain functional, even if the system has a complete power outage. For more information about NVDIMMs, see 2.6.5, “Backup Power Modules and NVDIMM” on page 65.

The iPDUs are managed by using the black and gray internal private networks. Each of the outlets can be individually monitored, and powered on or off. The iPDUs support Simple Network Management Protocol (SNMP), telnet, and a web interface.

DS8900F HMCs are responsible for system power control and monitoring by communicating to the network interfaces of the iPDUs and RPCs.

HMCs provide control and monitoring of the following items:

- ▶ AC power ON/OFF the whole system
- ▶ iPDU configuration
- ▶ iPDU health checking and error reporting
- ▶ Collecting power usage statistical data
- ▶ Controlling single power outlets during service actions
- ▶ iPDU firmware update

Figure 2-24 shows the Ethernet connections of the optional iPDUs for a DS8910F model 993. The ME provides the Ethernet ports for the iPDU Ethernet interfaces.




Figure 2-24 DS8910F model 993 iPDU Ethernet connections

Figure 2-25 shows the Ethernet connections of the iPDUs for a DS8910F model 994, which is equivalent to the iPDU configuration on a DS8980F or DS8950F without an expansion frame. All Ethernet ports of the ME private black and gray network are now occupied.




Figure 2-25 DS8910F model 994 iPDU Ethernet connections

Figure 2-26 on page 63 shows an example how the power is distributed when two iPDU pairs are installed. Note the power connections of a second HPFE Gen2 storage enclosure or the second I/O enclosure pair.

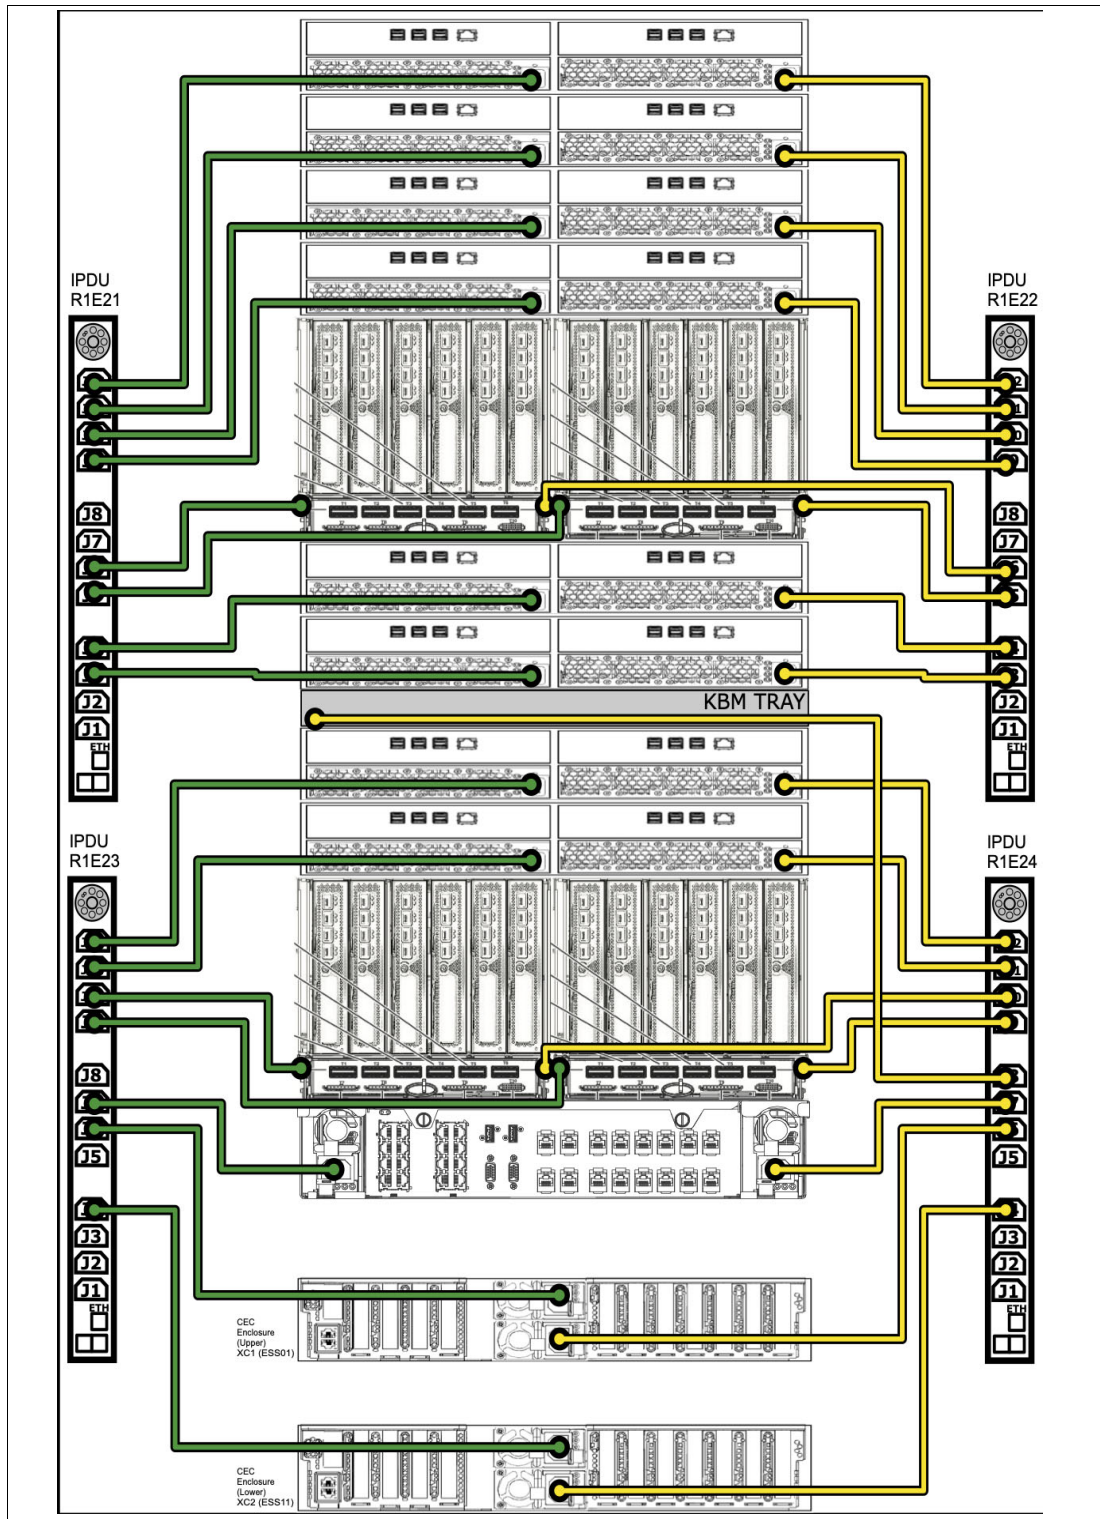


Figure 2-26 DS8910F model 994 iPDU power connections

Adding a model E96 expansion frame to a DS8980F or DS8950F system also adds another iPDU pair in that frame, which requires Ethernet connections to the internal management networks. To provide these connections, two extra Ethernet switches are installed in the base frame. This switch pair feature must be ordered with the expansion frame.

Figure 2-27 shows the two required network switches and the Ethernet connections for an expansion frame and three connected iPDU pairs.

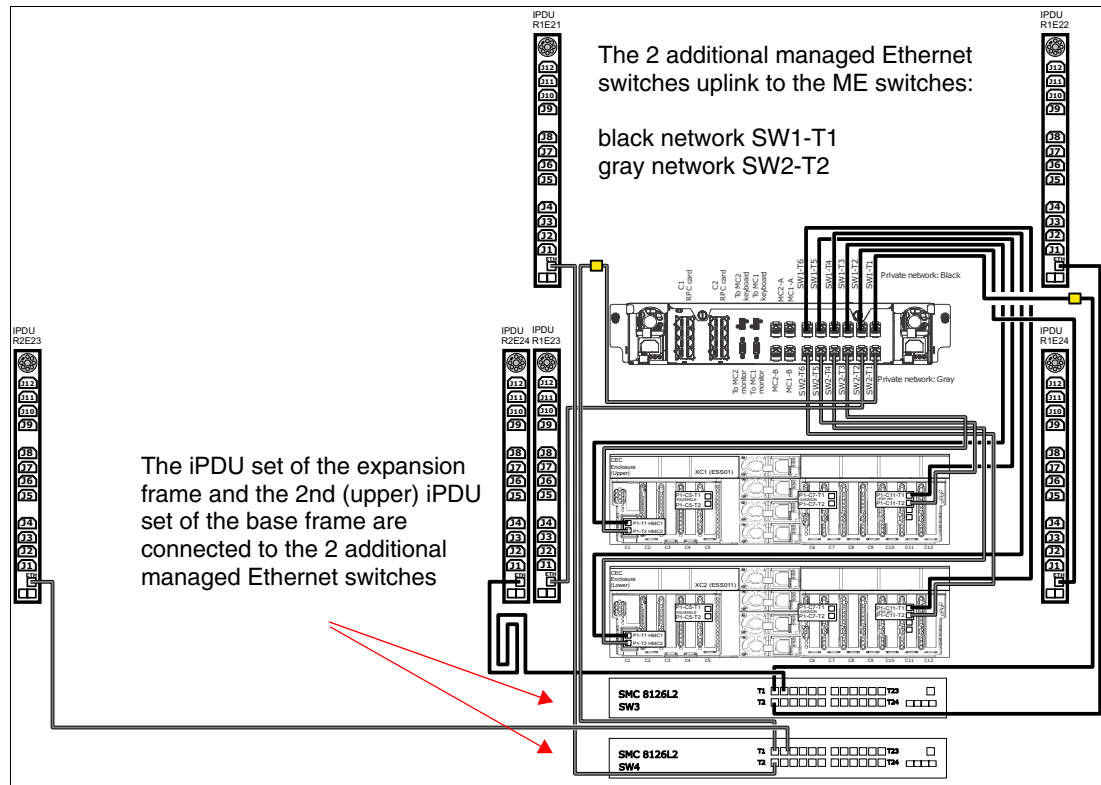


Figure 2-27 DS8950F model 996 and model E96 Ethernet connections

2.6.3 Power domains

To provide redundant power, the iPDUs are organized into groups that are called *power domains*. As shown in Figure 2-26 on page 63 and Figure 2-27, the iPDUs on one side of the racks form a power domain. If you look at the rear of the storage system, the iPDUs on the left side of the racks are in one domain (green), and the iPDUs on the right side of the racks are in the other (yellow) power domain.

The redundant power supplies in the CPCs, I/O enclosures, HPFE Gen2 enclosures, and the ME are connected across both power domains. The left power supplies connect to the green domain, while the right power supplies connect to the yellow domain.

For full redundancy, each power domain must be connected to separate power distribution systems that are fed by independent building power sources or service entrances.

2.6.4 Rack management 24-port Ethernet switch pair

When an expansion frame is ordered, an extra Ethernet switch pair is required to provide internal network connectivity to all installed iPDUs. The switches are installed at the bottom of the base frame. The two extra switches are called SW3 and SW4. One switch uplinks to the black and one to the gray network switch in the ME. For more information about connecting iPDUs for Ethernet management, see 2.6.2, “Intelligent Power Distribution Units” on page 59.

2.6.5 Backup Power Modules and NVDIMM

In DS8900F systems, system memory consists of RDIMMs for read/write cache memory, and NVDIMMs, which contain the write NVS partitions. A portion of the NVDIMM capacity is also accessed as ordinary read/write cache. The NVDIMM Persistent Memory Write Cache design eliminates the need for DC-UPSs and bulky battery sets, thus reducing the footprint of the DS8900 racks.

During normal operation, the NVDIMMs behave like any other DRAM, but when a power outage or other system failure occurs, the NVS partition contents are hardened in NAND flash storage. This NAND flash storage is with the DRAM chips on the NVDIMM module. The content is encrypted when written to the flash storage to prevent unauthorized access to the contents. Storing the write cache data in flash chips replaces the need for a *fire hose dump*, which was used on earlier DS8000 models to harden NVS data to disk. Figure 2-28 shows a symbolic view of an NVDIMM module.

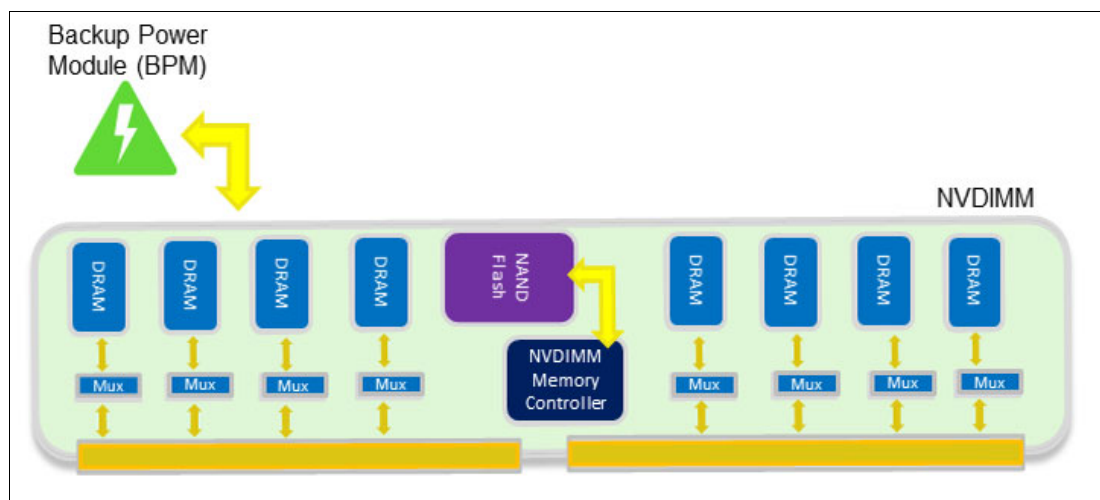


Figure 2-28 NVDIMM logic

BPMs connect directly to the NVDIMM modules to provide power during the DRAM to flash operation. They are specific nickel-based hybrid energy storage modules with a high-power discharge and fast charge times of 3 - 15 minutes. When system power is restored, the NVDIMMs move the preserved data from flash back to DRAM to be destaged to the storage system arrays during initial microcode load (IML).

The POWER9 processor-based systems support two NVDIMMs per CPC in designated memory slots.

The size of a BPM is smaller than a standard 2.5-inch disk drive module (DDM) and fits into one of the free CPC disk drive bays. A maximum of two BPMs are installed per CPC.

Figure 2-29 shows an example of a BPM.




Figure 2-29 Backup Power Module

With the BPMs connected directly to the NVDIMMs, the DRAM to flash operation functions independently without the need for any power that is provided by the CPC.

The NVDIMM capability is in addition to the data protection concept of storing the write cache NVS on the alternative node. For more information, see 3.2, “CPC failover and failback” on page 78.

NVDIMM configurations use either two 16 GB or two 32 GB modules. NVDIMMs are always installed in pairs of the same size. With 16 GB NVDIMMs, one BPM is sufficient to provide power to both modules in one CPC. With 32 GB NVDIMMs, two BPMs are provided in each CPC, with one for each NVDIMM.

2.6.6 Power cord options

The power cord must be ordered for a specific input voltage and connector type to meet local requirements, and these requirements vary worldwide. In cases where the supplied power cord does not include a connector, the proper connector must be installed by an electrician after the system is delivered. For more information, see the *IBM System Storage DS8900F Introduction and Planning Guide*, SC27-9560.

2.6.7 Enclosure power supply units

The CPCs, I/O enclosures, ME, and HPFE Gen2 enclosures have redundant PSUs for each enclosure that are fed by both left and right-side iPDUs. Left and right-side iPDUs should be fed by independent power distribution sources whenever possible. The PSUs have their own internal cooling fans. Each enclosure also has its own redundant cooling fans. All fans draw cool air from the front of the frame and exhaust hot air to the rear of the frame.

Note: The DS8900F is designed for efficient air flow and to be compliant with *hot and cold aisle* data center configurations.

2.7 Management Console and network

Every DS8900F base frame is equipped with a ME that includes two small form-factor (SFF) mini-PC HMCs, and two private network Ethernet switches. The secondary HMC is a redundant point of management in the DS8900F and it sits next to the primary HMC.

Figure 2-30 shows a diagram of the mini-PC HMCs and keyboard and monitor drawer location in the DS8950F model 996 base frame.




Figure 2-30 Diagram of mini-PC HMC and keyboard and monitor drawer location in the DS8950F

The storage administrator runs all DS8900F logical configuration tasks by using the Storage Management GUI or DS CLI. All client communications to the storage system are through the HMCs.

Clients that use the DS8900F advanced functions, such as MM or FlashCopy, communicate to the storage system with IBM Copy Services Manager.

The HMCs provide connectivity between the storage system and external Encryption Key Manager (EKM) servers.

HMCs also provide remote Call Home and remote support connectivity.

For more information about the HMC, see Chapter 6, “IBM DS8900F Management Console planning and setup” on page 167.

2.7.1 Ethernet switches

The DS8900F base frame has two 8-port Ethernet switches in the ME. The two switches provide redundant connectivity for the black and gray private management networks. The FSPs and LPARs in each CPC have dual network connections, with each connecting to the black and gray switches through the external breakout ports at the rear of the ME.

The switches receive power from the PSUs inside the ME and do not require separate power outlets. The ports on these switches are shown in Figure 2-31.




Figure 2-31 Eight-port Ethernet switches (SW1 and SW2) in the Management Enclosure

Each HMC also uses two designated Ethernet interfaces for the internal black (eth0) and gray (eth3) networks. Because the HMCs are installed in the ME, they are connected directly to the switches without routing through the external breakout ports.

The black and gray networks provide fully redundant communication between the HMCs and CPCs. These networks cannot be accessed externally, and no external connections are allowed. External customer network connections for both HMCs are provided at the rear of the base rack.

When an expansion frame is installed, the DS8900F has two 24-port switches (one each for the gray and black private networks) at the bottom of the base frame. These switches provide internal network connectivity to the iPDU in the expansion frame.

The 24-port switches are shown in Figure 2-32 on page 69.




Figure 2-32 Twenty-four-port Ethernet switches (SW3 and SW4) at the bottom of DS8950F model 996 for extra iPDUs

Important: The internal Ethernet switches that are shown in Figure 2-31 and Figure 2-32 are for the DS8900F private networks only. **Do not** connect an external network (or any other equipment) to the black or gray network switches.



IBM DS8900F reliability, availability, and serviceability

This chapter describes the reliability, availability, and serviceability (RAS) characteristics of the IBM DS8900F.

This chapter covers the following topics:

- ▶ DS8900F processor complex features
- ▶ CPC failover and failback
- ▶ Data flow in the DS8900F
- ▶ RAS on the Hardware Management Console
- ▶ RAS on the storage system
- ▶ RAS on the power subsystem
- ▶ Other features

3.1 DS8900F processor complex features

RAS is an important concept in the design of the DS8900F. Hardware features, software features, design considerations, and operational guidelines all contribute to make the DS8900F reliable. At the heart of the DS8900F is a pair of POWER9 processor-based servers. These servers, which are known as central processor complexes (CPCs), share the load of receiving and moving data between the attached hosts and the storage arrays. For more information, see *IBM Power Systems S922, S914, and S924 Technical Overview and Introduction*, REDP-5497.

However, the CPCs are also redundant so that if either one fails, the system switches to the remaining CPC and continues to run without any host I/O interruption. This section looks at the RAS features of the CPCs, including the hardware, the operating system (OS), and the interconnections.

3.1.1 POWER9 PowerVM Hypervisor

The POWER9 IBM PowerVM® Hypervisor (PHYP) is a component of system firmware that is always active regardless of the system configuration, even when it is not connected to the Management Console (MC). PHYP runs on the flexible service processor (FSP). It requires the FSP processor and memory to support the resource assignments to the logical partition (LPAR) on the server. It operates as a hidden partition with no CPC processor resources that are assigned to it, but it does allocate a small amount of memory from the partition.

The PHYP provides the following capabilities:

- ▶ Reserved memory partitions set aside a portion of memory to use as cache and a portion to use as non-volatile storage (NVS) functioning as write cache.
- ▶ Preserved memory support allows the contents of the NVS and cache areas to be protected if a server restarts.
- ▶ I/O enclosure initialization, power control, and slot power control prevent a CPC that is restarting from initializing an I/O adapter that is in use by another server.
- ▶ It provides automatic restart of a hung or stopped partition. The PHYP also monitors the service processor and runs a reset or reload if it detects the loss of the service processor. It notifies the OS if the problem is not corrected.

The AIX OS uses PHYP services to manage the Translation Control Entry (TCE) tables. The OS communicates the wanted I/O bus address to logical mapping, and the PHYP returns the I/O bus address to physical mapping within the specific TCE table. The PHYP needs a dedicated memory region for the TCE tables to convert the I/O address to the partition memory address. The PHYP then can monitor direct memory access (DMA) transfers to the Peripheral Component Interconnect Express (PCIe) adapters.

3.1.2 POWER9 processor

The POWER9 processor implements 64-bit IBM Power Architecture® technology. The multi-core architecture of the POWER9 processor-based modules is matched with innovation across a wide range of related technologies to deliver leading throughput, efficiency, scalability, and RAS.

Areas of innovation, enhancement, and consolidation

The POWER9 processor represents an important performance increase that can be as high as 70% over previous comparable hardware generations. The POWER9 processor has the following areas of innovation, enhancement, and consolidation:

- ▶ A large on-chip, low-latency L3 cache that is implemented in embedded dynamic random access memory (eDRAM). Lower energy consumption and a smaller physical footprint are two of the benefits.
- ▶ Cache hierarchy and component innovation.
- ▶ Advances in memory subsystem with eight memory channels.
- ▶ Advances in off-chip signaling.
- ▶ The POWER9 processor has intelligent threads that can vary based on the workload demand. The system automatically selects whether a workload benefits from dedicating as much capability as possible to a single thread of work, or if the workload benefits more from spreading the capability across up to 44 threads (four per core) of work by using simultaneous multithreading (SMT). With more threads, the POWER9 processor can deliver more total capacity as more tasks are accomplished in parallel. With fewer threads, those workloads that need fast individual tasks can get the performance that they need for maximum benefit.

The remainder of this section describes the RAS features of the POWER9 processor. These features and abilities apply to the DS8900F. You can read more about the POWER9 and processor configuration from the DS8900F architecture point of view in 2.3.1, “IBM POWER9 processor-based CPCs” on page 42.

POWER9 RAS features

The following sections describe the RAS leadership features of IBM POWER9 processor-based systems.

POWER9 processor instruction retry and recovery (IRR)

As with previous generations, the POWER9 processor can run processor instruction retry and alternative processor recovery for many core-related faults. This ability reduces exposure to permanent and intermittent errors in the processor core.

With the instruction retry function, when an error is encountered in the core in caches and certain logic functions, the POWER9 processor first automatically retries the instruction. If the source of the error was truly transient, the instruction succeeds and the system can continue normal operation.

POWER9 cache protection and cache error handling

The processor instruction retry function protects processors and data caches. The L1 cache is divided into sets. The POWER9 processor can deallocate all but one set before a processor instruction retry is run. In addition, faults in the segment lookaside buffer (SLB) array are recoverable by the PHYP. The SLB is used in the core to run address translation calculations.

The L2 and L3 caches in the POWER9 processor are protected with double-bit detect single-bit correct error correction code (ECC). Single-bit errors are corrected before they are forwarded to the processor, and then they are written back to L2 or L3.

In addition, the caches maintain a cache line delete capability. A threshold of correctable errors that is detected on a cache line can result in purging the data in the cache line and removing the cache line from further operation without requiring a restart. An ECC uncorrectable error that is detected in the cache can also trigger a purge and delete of the cache line.

This action results in no loss of operation because an unmodified copy of the data can be held in system memory to reload the cache line from main memory. Modified data is handled through special uncorrectable error handling. L2 and L3 deleted cache lines are marked for persistent deconfiguration on subsequent system restarts until they can be replaced.

POWER9 first-failure data capture

First-failure data capture (FFDC) is an error isolation technique. FFDC ensures that when a fault is detected in a system through error checkers or other types of detection methods, the root cause of the fault is captured without the need to re-create the problem or run an extended tracing or diagnostics program.

For most faults, a good FFDC design means that the root cause is detected automatically without intervention by an IBM Systems Service Representative (IBM SSR). Pertinent error data that relates to the fault is captured and saved for analysis. In hardware, FFDC data is collected from the fault isolation registers and the associated logic. In firmware, this data consists of return codes, function calls, and other items.

FFDC *check stations* are carefully positioned within the server logic and data paths to ensure that potential errors can be identified quickly and accurately tracked to a field-replaceable unit (FRU).

This proactive diagnostic strategy is an improvement over the classic, less accurate *restart and diagnose* service approach.

Redundant components

High opportunity components (those components that most affect system availability) are protected with redundancy and the ability to be repaired concurrently.

The following redundant components allow the system to remain operational:

- ▶ POWER9 cores, which include redundant bits in L1 instruction and data caches, L2 caches, and L2 and L3 directories.
- ▶ IBM Power S922 and IBM Power S924 use CPC main memory with Direct Attach Industry Standard DIMMs (ISDIMMs), which use an ECC algorithm that improves single-bit error correction and memory failure identification.
- ▶ Redundant cooling.
- ▶ Redundant power supply units (PSUs).
- ▶ Redundant links to the I/O subsystem.
- ▶ Concurrent maintenance for PCI adapters.

Self-healing

For a system to be self-healing, it must be able to recover from a failing component by detecting and isolating the failed component. The system is then able to take the component offline, fix, or isolate it, and then reintroduce the fixed or replaced component into service without any application disruption. Self-healing technology includes the following examples:

- ▶ *Chipkill*, which is an enhancement that enables a system to sustain the failure of an entire DRAM chip. The system can continue indefinitely in this state with no performance degradation until the failed dual inline memory module (DIMM) can be replaced.
- ▶ Single-bit error correction by using ECC without reaching error thresholds for main, L2, and L3 cache memory.
- ▶ L2 and L3 cache line delete capability, which provides more self-healing.

- ▶ ECC extended to inter-chip connections on the fabric and processor bus.
- ▶ Dynamic processor deallocation.

Bus cyclic redundancy check and lane repair

ECC is used internally in various data paths as data is transmitted between units. High-speed data buses can be susceptible to the occasional multiple bit errors due to the nature of the bus design. A cyclic redundancy check (CRC) code is used to determine whether there are errors within an entire packet of data. If a bit error is recognized, the bus can retrain and retry the operation and continue. CRC checking is done for the memory bus, and in POWER9 CRC checking is now done for the processor fabric bus interfaces sending data between processors.

The memory bus between processors and the memory uses CRC with retry. The design also includes a spare data lane so that if a persistent single data error exists, the faulty bit can be “self-healed.” The POWER9 busses between processors also have a spare data lane that can be substituted for a failing one to “self-heal” the single bit errors.

Memory reliability, fault tolerance, and integrity

POWER9 uses ECC circuitry for system memory to correct single-bit memory failures. The ECC algorithm works on ISDIMM pairs on a rank basis. With this ECC code, the system can dynamically allow correction from an entire DRAM failure (Chipkill) correction by using x4 DRAMs. It can also correct an error even if another symbol (a byte, which is accessed by a 2-bit line pair) experiences a fault.

A rank of four ISDIMMs contains enough DRAMs to provide 64 bits of data at a time with enough check bits to correct the case of a single DRAM module after the bad DRAM is detected, and then correct an extra faulty bit.

The ability to correct an entire DRAM is what IBM traditionally called *Chipkill correction*. Correcting this kind of fault is essential in protecting against a memory outage and should be considered as a minimum error correction for any modern server design.

The POWER9 processors that are used in DS8900F are designed internally for ISDIMMs without an external buffer chip. The ECC checking is at the 64-bit level, so Chipkill protection is provided with x4 DIMMs plus some additional sub Chipkill level error checking after a Chipkill event.

The memory DIMMs also use hardware scrubbing and thresholding to determine when memory modules within each bank of memory must be used to replace modules that exceeded their threshold of error count. Hardware scrubbing is the process of reading the contents of the memory during idle time and checking and correcting any single-bit errors that accumulated by passing the data through the ECC logic. This function is a hardware function on the memory controller chip, and does not influence normal system memory performance.

The ability to use hardware accelerated scrubbing to refresh memory that might have experienced soft errors is a given. The memory bus interface is also important. The direct bus-attach memory that is used in the scale-out servers supports RAS features in that design, including register clock driver (RCD) parity error detection and retry.

Fault masking

If corrections and retries succeed and do not exceed threshold limits, the system remains operational with full resources, and no external administrative intervention is required.

Mutual surveillance

The service processor monitors the operation of the IBM POWER Hypervisor firmware during the boot process and monitors for loss of control during system operation. It also allows the POWER Hypervisor to monitor service processor activity. The service processor can take the correct action (including calling for service) when it detects that the POWER Hypervisor firmware lost control. The POWER Hypervisor can also request a service processor repair action if necessary.

3.1.3 Cross-cluster communication

In the DS8900F, the I/O enclosures are connected point-to-point, and each CPC uses a PCIe architecture. DS8900F uses the PCIe paths between the I/O enclosures to provide the cross-cluster (XC) communication between CPCs. This configuration means that no separate path is between the XC communications and I/O traffic, which simplifies the topology. During normal operations, the XC communication traffic uses a small portion of the overall available PCIe bandwidth (less than 1.7%), so XC communication traffic has a negligible effect on I/O performance.

Figure 3-1 shows the redundant PCIe fabric design for XC communication in the DS8900F and depicts the single-chip modules (SCMs) (SCM #0 and SCM#1) in each CPC. If the I/O enclosure that is used as the XC communication path fails, the system automatically uses an available alternative I/O enclosure for XC communication.

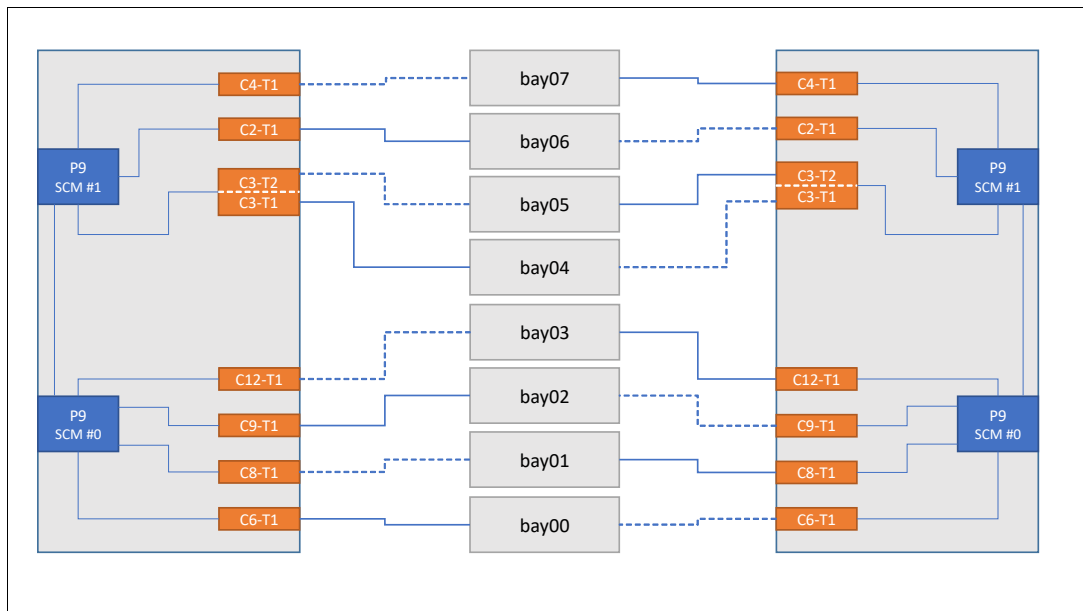


Figure 3-1 DS8900F XC communication through the PCIe fabric and I/O enclosures

3.1.4 Environmental monitoring

The environment (power, fans, and temperature) is monitored by the FSP. Environmental critical and non-critical conditions generate emergency power-off warning (EPOW) events. Critical events (for example, a complete input power loss) trigger the correct signals from the hardware to start an emergency shutdown to prevent data loss without OS or firmware involvement. Non-critical environmental events are logged and reported by using Event Scan.

The temperature is also monitored. If the ambient temperature rises above a preset operating range, the rotational speed of the cooling fans increases. Temperature monitoring also warns the Licensed Internal Code (LIC) of potential environmental problems. An orderly system shutdown, including a service call to IBM, occurs when the operating temperature exceeds a critical level.

Voltage monitoring provides a warning and an orderly system shutdown when the voltage is out of the operational specification range.

More monitoring support can be found by running the DS CLI **showsu** command and viewing the Added Energy Report (ER) Test Mode, ER Recorded, ER Power Usage, ER Inlet Temp, ER I/O Usage, and ER Data Usage fields, as shown in Example 3-1.

Example 3-1 The showsu command

```
dsccli> showsu
Name           -
desc           Sand Shark
ID             IBM.2107-75HAL90
Model         996
WWNN          5005076309FFEC62
config        Undefined
pw state      On
pw mode       Remote Manual
reqpm        Remote Manual
System Memory 512.0 GB
MTS           IBM.5331-75HAL90
ER Test Mode Disabled
ER Recorded 2022-06-22T22:39:03+0200
ER Power Usage 1730
ER Inlet Temp 23.0
ER I/O Usage 83
ER Data Usage 458
```

3.1.5 Resource deconfiguration

If recoverable errors exceed threshold limits, resources can be unconfigured to keep the system operational. This ability allows deferred maintenance at a more convenient time. Dynamic deconfiguration of potentially failing components is nondisruptive, which allows the system to continue to run. Persistent deconfiguration occurs when a failed component is detected. It is then deactivated on a subsequent restart.

Dynamic deconfiguration functions include the following components:

- ▶ Processor
- ▶ L3 cache lines
- ▶ Partial L2 cache deconfiguration
- ▶ PCIe bus and slots

Persistent deconfiguration functions include the following components:

- ▶ Processor
- ▶ Memory
- ▶ Unconfigure or bypass failing I/O adapters
- ▶ L2 cache

After a hardware error is flagged by the service processor, the subsequent restart of the CPC starts extended diagnostic testing. If a processor or memory is marked for persistent deconfiguration, the boot process attempts to proceed to completion with the faulty device automatically unconfigured. Failing I/O adapters are unconfigured or bypassed during the boot process.

3.2 CPC failover and failback

To understand the process of CPC failover and failback, you must review the logical architecture of the DS8900F. For more information, see Chapter 4, “Virtualization concepts” on page 107.

3.2.1 Dual cluster operation and data protection

For processing host data, a basic premise of RAS is that the DS8000 system always tries to maintain two copies of write data while the data moves through the storage system. Two areas of the primary memory of the nodes are used for holding host data: cache memory and NVS.

For a DS8980F model 998 with 4.3 TB total system memory or a DS8950F system with a maximum configuration of 3.4 TB of total system memory, NVS is 128 GB. For IBM DS8910F model 993 and DS8910F model 994, all configurations use 1/16th of system memory except for the smallest systems with 192 GB of total system memory, which uses the minimum of 8 GB of NVS. NVS contains write data until the data is destaged from cache to the drives. NVS data is protected and kept by non-volatile dual inline memory module (NVDIMM) technology, where the data is moved from DRAM to a flash memory on the NVDIMM modules if the DS8900F experiences a complete loss of input AC power.

When a write is sent to a volume and both the nodes are operational, the *write data* is placed into the cache memory of the owning node and into the NVS of the other CPC. The NVS copy of the write data is accessed only if a write failure occurs and the cache memory is empty or possibly invalid. Otherwise, the NVS copy of the write data is discarded after the destaging from cache to the drives is complete.

The location of write data when both CPCs are operational is shown in Figure 3-2 on page 79, which shows how the cache memory of node 0 in CPC0 is used for all logical volumes that are members of the even logical subsystems (LSSs). Likewise, the cache memory of node 1 in CPC1 supports all logical volumes that are members of odd LSSs. For every write that is placed into cache, a copy is placed into the NVS memory that is in the alternative node. Therefore, the following normal flow of data for a write when both CPCs are operational is used:

1. Data is written to cache memory in the owning node. At the same time, data is written to the NVS memory of the alternative node.
2. The write operation is reported to the attached host as complete.
3. The write data is destaged from the cache memory to a drive array.
4. The write data is discarded from the NVS memory of the alternative node.

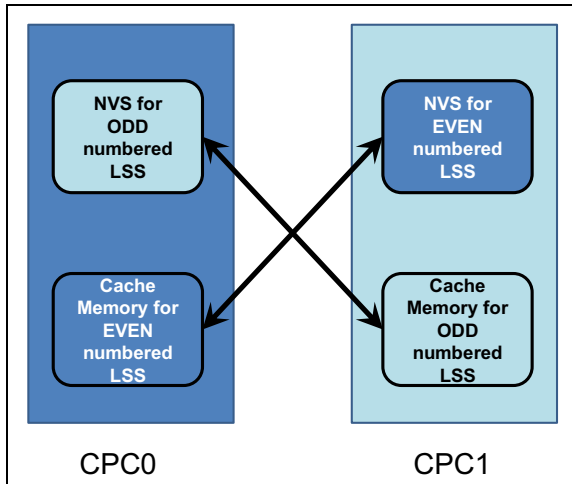


Figure 3-2 NVS write data when both CPCs are operational

Under normal operation, both DS8900F nodes are actively processing I/O requests. The following sections describe the failover and failback procedures that occur between the CPCs when an abnormal condition affects one of them.

3.2.2 Failover

In the example that is shown in Figure 3-3, CPC0 failed. CPC1 must take over all of the CPC0 functions. All storage arrays are accessible by both CPCs.

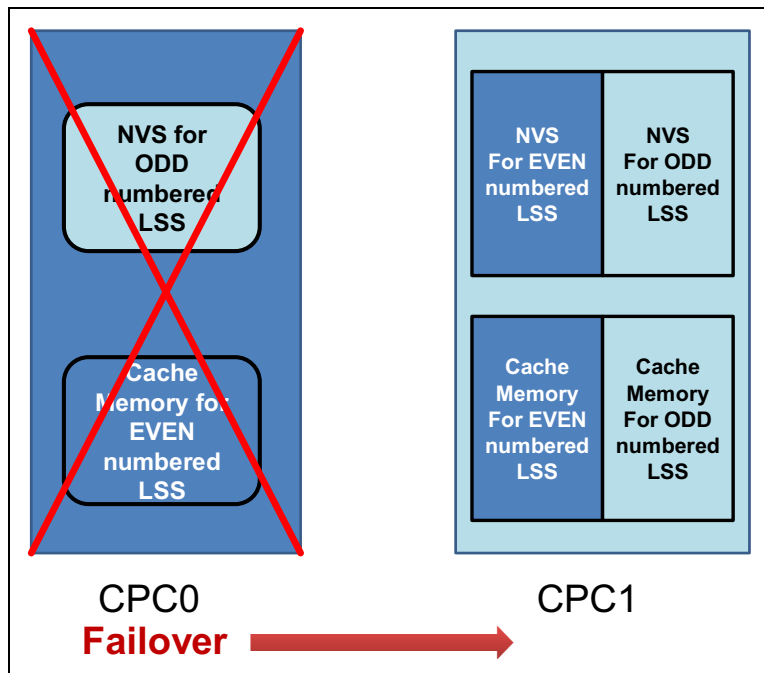


Figure 3-3 CPC0 failover to CPC1

At the moment of failure, node 1 in CPC1 includes a backup copy of the node 0 write data in its own NVS. From a data integrity perspective, the concern is the backup copy of the node 1 write data, which was in the NVS of node 0 in CPC0 when it failed. Because the DS8900F now has only one copy of that data (active in the cache memory of node 1 in CPC1), it performs the following steps:

1. Node 1 destages the contents of its NVS (the node 0 write data) to the drive subsystem. However, before the actual destage and at the beginning of the failover, the following tasks occur:
 - a. The surviving node starts by preserving the write data in cache that was backed up by the failed CPC NVS. If a restart of the single working CPC occurs before the cache data is destaged, the write data remains available for subsequent destaging.
 - b. The existing write data in cache (for which only a single volatile copy exists) is added to the NVS so that it remains available if the attempt to destage fails or a server restart occurs. This function is limited so that it cannot use more than 85% of NVS space.
2. The NVS and cache of node 1 are divided in two portions, one for the odd LSSs and one for the even LSSs.
3. Node 1 begins processing the I/O for all the LSSs, taking over for node 0.

This entire process is known as a *failover*. After failover, the DS8900F operates as shown in Figure 3-3 on page 79. Node 1 now owns all the LSSs, which means all reads and writes are serviced by node 1. The NVS inside node 1 is now used for both odd and even LSSs. The entire failover process is transparent to the attached hosts.

The DS8900F can continue to operate in this state indefinitely. No functions are lost, but the redundancy is lost, and performance is decreased because of the reduced system cache. Any critical failure in the working CPC renders the DS8900F unable to serve I/O for the arrays, so the IBM Support team begins work immediately to determine the scope of the failure and build an action plan to restore the failed CPC to an operational state.

3.2.3 Failback

The *failback* process begins automatically when the DS8900F determines that the failed CPC did not resume an operational state. If the failure was relatively minor and recoverable by the DS8900F OS, the software starts the resume action. If a service action occurred and hardware components were replaced, the IBM SSR or remote support engineer resumes the failed CPC.

This example in which CPC0 failed assumes that CPC0 was repaired and resumed. The failback begins with server 1 in CPC1 starting to use the NVS in node 0 in CPC0 again, and the ownership of the even LSSs being transferred back to node 0. Normal I/O processing, with both CPCs operational, then resumes. Just like the failover process, the failback process is transparent to the attached hosts.

In general, recovery actions (failover or failback) on the DS8900F do not affect I/O operation latency by more than 8 seconds.

If you require real-time response in this area, contact IBM to determine the latest information about how to manage your storage to meet your requirements.

3.2.4 NVS and power outages

DS8900F systems contain up to two pairs of intelligent power distribution units (iPDUs) in the base frame, and a third pair of iPDUs when a model E96 expansion frame is installed. One iPDU in each pair is in the green power domain, and its partner is in the yellow power domain.

During normal operation, the DS8900F preserves write data by storing a duplicate copy in the NVS of the alternative CPC. To ensure that write data is not lost during a power failure event, the DS8900F stores the NVS contents on non-volatile DIMMs (NVDIMMs). Each CPC contains two NVDIMMs with dedicated Backup Power Modules (BPMs). The NVDIMMs act as regular DRAM during normal operation. During AC power loss, the BPMs provide power to the NVDIMM modules until they have moved all modified data (NVS) to integrated flash memory. The NVDIMM save process is autonomous, and requires nothing from the CPC.

Important: DS8900F can tolerate a power line disturbance (PLD) for up to 20 ms. A PLD that exceeds 20 ms on both power domains initiates an emergency shutdown.

The following sections describe the steps that occur when AC input power is lost to both power domains.

Power loss

When a wall power loss condition occurs, the following events occur:

1. All host adapter I/O is blocked.
2. Each NVDIMM begins copying its NVS data to the internal flash partition.
3. The system powers off without waiting for the NVDIMM copy operation.
4. The copy process continues and completes independently from the storage systems power.

Power restored

When power is restored, the DS8900F must be powered on manually unless the remote power control mode is set to *automatic*.

Note: Be careful if you decide to set the remote power control mode to *automatic*. If the remote power control mode is set to *automatic*, after input power is restored, the DS8900F is powered on automatically.

For more information about how to set power control on the DS8900F system, see [IBM Documentation](#).

After the DS8900F is powered on, the following events occur:

1. The CPCs are powered on, PHYP loads, and power-on self-test (POST) runs.
2. Each CPC boots up, and begins the initial microcode load (IML).
3. At an early stage in the IML process, the CPC detects NVS data on its NVDIMMs and restores the data to destage it to the storage drives.

3.3 Data flow in the DS8900F

The DS8900F connectivity between the CPC and the I/O enclosures uses the PCIe architecture. For more information, see 2.3.5, “Peripheral Component Interconnect Express adapters” on page 47.

3.3.1 I/O enclosures

As shown in Figure 3-1 on page 76, each CPC on a DS8950F is connected to two or four I/O enclosures in the base frame, and up to eight I/O enclosures when the expansion frame is installed. Each I/O enclosure functions as an extension the CPCs. The DS8910F model 993 has a maximum of two I/O enclosures. The DS8910F model 994 has base of two I/O enclosures and a maximum of four. The DS8980F Model 998 and DS8950F model 996 start with two I/O enclosures up to a maximum of eight.

The DS8900F I/O enclosures use adapters with PCIe connections. The adapters in the I/O enclosures are concurrently replaceable. Each slot can be independently powered off for installation, replacement, or removal of an adapter.

In addition, each I/O enclosure has $N+1$ power and cooling redundancy in the form of two PSUs with integrated fans, and two enclosure cooling fans. The PSUs and enclosure fans can be replaced concurrently without disruption to the I/O enclosure.

3.3.2 Host connections

Each DS8900F 32 Gbps or 16 Gbps Fibre Channel (FC) host adapter provides four longwave (LW) or shortwave (SW) ports for connectivity to storage area network (SAN) switches, or directly to hosts. Each port can be independently configured for FCP or FICON topology.

Single or multiple paths

The host adapters are shared between the CPCs. To illustrate this concept, Figure 3-4 on page 83 shows a potential system configuration. In this example, two I/O enclosures are shown. Each I/O enclosure has up to four FC host adapters. If a host server has only a single path to a DS8900F, as shown in Figure 3-4 on page 83, it can access volumes that belong to all LSSs because the host adapter directs the I/O to the correct CPC. However, if an error occurs on the host adapter, host adapter port, I/O enclosure, or in the SAN, all connectivity is lost because this configuration has no redundancy. The same is true for the host bus adapter (HBA) in the attached host, making it a single point of failure (SPoF) without a redundant HBA.

Important: For host connectivity, hosts that access the DS8900F must have at least two connections to I/O ports on separate host adapters in separate I/O enclosures.

Figure 3-4 shows a single-path host connection.

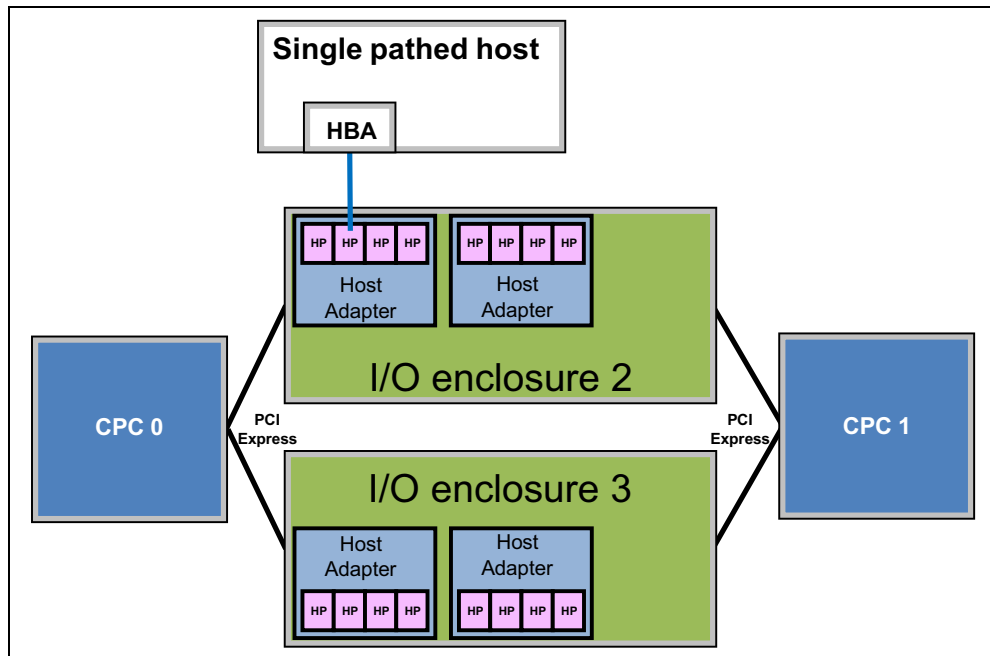


Figure 3-4 A single-path host connection

A more robust design is shown in Figure 3-5, in which the host is attached to separate FC host adapters in separate I/O enclosures. This configuration is also important because during a LIC update, a host adapter port might need to be taken offline. This configuration allows host I/O to survive a hardware failure on any component on either path.

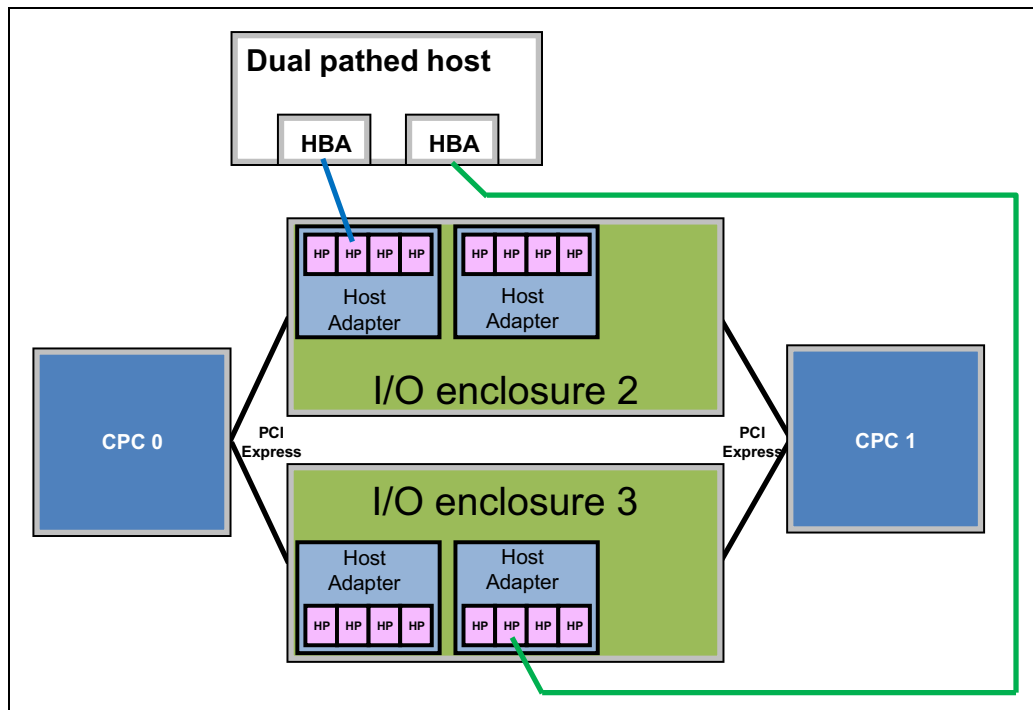


Figure 3-5 A dual-path host connection

SAN and FICON switches

Because many hosts can connect to the DS8900F by using multiple paths, the number of host adapter ports that are available in the DS8900F might not be sufficient to accommodate all the connections. The solution to this problem is to use SAN switches or directors to switch logical connections from multiple hosts. In an IBM Z environment, a SAN switch or director that supports FICON is required.

A logic or power failure in a switch or director can interrupt communication between hosts and the DS8900F. Provide more than one switch or director to ensure continued availability. Configure ports from two separate host adapters in two separate I/O enclosures to go through each of two directors. The complete failure of either director leaves the paths that are configured to the alternative director still available.

Support for the T10 Data Integrity Field standard

The DS8900F incorporates the American National Standards Institute (ANSI) T10 Data Integrity Field (DIF) standard for Fixed-Block (FB) volumes.

When data is read, the DIF is checked before the data leaves the DS8900F and again when the data is received by the host system. Previously, it was possible to ensure the data integrity within the storage system only with ECC. However, T10 DIF can now check end-to-end data integrity through the SAN. Checking is done by hardware, so no performance impact occurs.

For more information about T10 DIF implementation in the DS8900F, see “T10 Data Integrity Field support” on page 118.

Robust and resilient SAN data transfer

Forward Error Correction (FEC) is enabled on 16 Gbps and 32 Gbps host adapters to intercept and correct bit errors for utmost reliability. FEC is also implemented on IBM z13 and later systems, providing end-to-end reliability.

To provide more proactive system diagnosis information about SAN fabric systems, the *Read Diagnostic Parameters* (RDP) function, which complies with industry standards, is implemented on the DS8900F. This function provides host software with the capability to perform predictive failure analysis (PFA) on degraded SAN links before they fail.

When troubleshooting SAN errors, the IBM SSR can run a wrap test on a single host adapter port without taking the entire adapter offline.

Multipathing software

Each attached host OS requires multipathing software to manage multiple paths to the same device, and to provide redundant routes for host I/O requests. When a failure occurs on one path to a logical device, the multipathing software on the attached host can identify the failed path and route the I/O requests for the logical device to alternative paths. Furthermore, it can likely detect when the path is restored. The multipathing software that is used varies by attached host OS and environment, as described in the following sections.

Open systems

In most open systems environments, multipathing is available at the OS level. The Subsystem Device Driver (SDD), which was provided and maintained by IBM for several OSs, is now an obsolete approach for a multipathing solution.

Important: The IBM System Storage Multipath Subsystem Device Driver is no longer supported for IBM DS8000. The Subsystem Device Driver Device Specific Module (SDDDSM) for Windows and Subsystem Device Driver Path Control Module (SDDPCM) for AIX are end of service (EOS). Clients must move to the native multipath OS software to get continuous support from IBM. For more information, see [SDDDSM and SDDPCM End Of Support for DS8000](#).

For the AIX OS, the DS8000 is supported through the AIX multipath I/O (MPIO) framework, which is included in the base AIX OS. Use the base AIX Multipath I/O Path Control Module (AIXPCM) support instead of the old SDDPCM.

For multipathing under Microsoft Windows, the DS8000 is supported by the native Microsoft MPIO stack by using Microsoft Device Specific Module (MSDSM). Existing environments that rely on the old SDDDSM should be moved to the native OS driver.

Note: To move existing SDDPCM and SDDDSM implementations, see the following resources:

- ▶ [How To Migrate SDDPCM to AIXPCM](#)
- ▶ [Migrating from SDDDSM to Microsoft MSDSM - SVC/Storwize](#)

For all newer versions of RHEL and SUSE Linux Enterprise Server, the native Linux multipathing driver, Device-Mapper Multipath (DM Multipath), is used.

Also, on the VMware vSphere ESXi server, the VMware Native Multipathing Plug-in (NMP) is the supported multipathing solution.

For more information about the multipathing software that might be required for various OSs, see the [IBM System Storage Interoperation Center \(SSIC\)](#).

IBM Z

In the IBM Z environment, a best practice is to provide multiple paths from each host to a storage system. Typically, four or eight paths are configured. The channels in each host that can access each logical control unit (LCU) in the DS8900F are defined in the hardware configuration definition (HCD) or input/output configuration data set (IOCDS) for that host. Dynamic Path Selection (DPS) allows the channel subsystem to select any available (non-busy) path to start an operation to the disk subsystem. Dynamic Path Reconnect (DPR) allows the DS8900F to select any available path to a host to reconnect and resume a disconnected operation, for example, to transfer data after disconnection because of a cache miss.

These functions are part of the IBM z/Architecture®, and are managed by the channel subsystem on the host and the DS8900F.

A physical FICON path is established when the DS8900F port sees light on the fiber, for example, a cable is plugged in to a DS8900F host adapter, a processor or the DS8900F is powered on, or a path is configured online by z/OS. Logical paths are established through the port between the host, and part or all of the LCUs in the DS8900F are controlled by the HCD definition for that host. This configuration happens for each physical path between an IBM Z host and the DS8900F. Multiple system images can be in a CPU. Logical paths are established for each system image. The DS8900F then knows the paths that can be used to communicate between each LCU and each host.

Control-unit initiated reconfiguration (CUIR) varies off a path or paths to all IBM Z hosts to allow service to an I/O enclosure or host adapter, then varies on the paths to all host systems when the host adapter ports are available. This function automates channel path management in IBM Z environments in support of selected DS8900F service actions.

CUIR is available for the DS8900F when it operates in the z/OS and IBM z/VM® environments. CUIR provides automatic channel path vary on and vary off actions to minimize manual operator intervention during selected DS8900F service actions.

CUIR also allows the DS8900F to request that all attached system images set all paths that are required for a particular service action to the offline state. System images with the correct level of software support respond to such requests by varying off the affected paths, and either notifying the DS8900F system that the paths are offline, or that it cannot take the paths offline. CUIR reduces manual operator intervention and the possibility of human error during maintenance actions, and reduces the time that is required for the maintenance. This function is useful in environments in which many z/OS or z/VM systems are attached to a DS8900F.

3.3.3 Metadata checks

When application data enters the DS8900F system, special codes or *metadata*, also known as *redundancy checks*, are appended to that data. This metadata remains associated with the application data while it is transferred throughout the DS8900F. The metadata is checked by various internal components to validate the integrity of the data as the data moves throughout the disk system. It is also checked by the DS8900F before the data is sent to the host in response to a read I/O request. The metadata also contains information that is used as an extra level of verification to confirm that the data that is returned to the host is coming from the location that you want on the disk.

The metadata check is independent of the DS8900F T10 DIF support for FB volumes. For more information about T10 DIF implementation in the DS8000, see “T10 Data Integrity Field support” on page 118.

3.4 RAS on the Hardware Management Console

The Hardware Management Console (HMC) is used to configure, manage, and maintain the DS8900F. Two HMCs (the *primary* and the *secondary*) are included in every DS8900F Management Enclosure (ME) in the base frame. The DS8900F HMCs work with IPv4, IPv6, or a combination of both IP standards. For more information about the HMC and network connections, see 6.1.1, “Management Enclosure” on page 168 and 5.3, “Network connectivity planning” on page 159.

The HMC is the DS8900F management focal point. If no HMC is operational, it is impossible to run maintenance, modifications to the logical configuration, or Copy Services (CS) tasks, such as the establishment of FlashCopy backups, Metro Mirror (MM) or Global Mirror (GM), by using the DS Command-line Interface (DS CLI), Storage Management GUI, or IBM Copy Services Manager. The implementation of a secondary HMC provides a redundant management focal point and is especially important if CS or Encryption Key Manager (EKM) are used.

3.4.1 Licensed Internal Code updates

The DS8900F contains many discrete redundant components. The DS8900F architecture allows concurrent code updates. This ability is achieved by using the redundant design of the DS8900F. The following components have firmware that can be updated concurrently:

- ▶ FSP and IBM Power firmware
- ▶ iPDU
- ▶ Rack power control cards (RPCCs)
- ▶ Host adapters
- ▶ Flash enclosure
- ▶ Device adapters (DAs)
- ▶ Flash drives
- ▶ I/O enclosure

The DS8900F CPCs have an OS (AIX) and Licensed Machine Code (LMC) that can be updated. As IBM continues to develop and improve the DS8900F, new releases of firmware and LMC become available that offer improvements in function and reliability. For more information about LIC updates, see Chapter 11, “Licensed Machine Code” on page 405.

3.4.2 Call Home and remote support

This section describes the Call Home feature and remote support capability.

Call Home

Call Home is the capability of the DS8900F to notify the client and IBM Support to report a problem. Call Home is configured in the HMC at installation time. Call Home to IBM Support is done over the customer network through a secure protocol. Customer notifications can also be configured as email (SMTP) or Simple Network Management Protocol (SNMP) alerts. An example of an email notification output is shown in Example 3-2.

Example 3-2 Typical email notification output

```
REPORTING SF MTMS: 2107-996*75HAL90
FAILING SF MTMS: 5331-996*75HAL90
REPORTING SF LPAR: unknown
PROBLEM NUMBER: 270
PROBLEM TIMESTAMP: Oct 16, 2019 3:15:46 PM CEST
REFERENCE CODE: BE83CB93

***** START OF NOTE LOG *****
BASE RACK ORDERED MTMS 5331-996*75HAL90
LOCAL HMC MTMS 8100LI7*I336057 HMC ROLE Primary
LOCAL HMC OUTBOUND CONFIG SSL only FTP: enabled
REMOTE HMC MTMS 8100LI7*I336145 HMC ROLE Secondary
REMOTE HMC OUTBOUND CONFIG Internet only FTP: enabled
AOS STATUS Running AOS VERSION 4.0
AOS ACL=(DS8k, Storage) AOS TRACE Enable
RSC STATUS Running RSC PASSWORD Required
HMC CE default HMC REMOTE default
HMC PE default HMC DEVELOPER default
2107 BUNDLE 89.0.208.0
HMC BUILD 1909192353
LMC LEVEL v25.90.0 build level 20191013.1
FIRMWARE LEVEL SRV0 01VL93087 SRV1 01VL93087

STORAGE FACILITY 2107-996*75HAL90,
PARTITION NAME SF75HAL90ESS11
```

PARTITION HOST NAME SF75HAL90ESS11
PARTITION STATUS SFI 2107-996*75HAL91 SVR 9009-42A*785A6C0 LPAR SF75HAL90ESS11 STATE =
AVAILABLE

FIRST REPORTED TIME Oct 16, 2019 3:15:46 PM CEST
LAST REPORTED TIME Oct 16, 2019 3:15:46 PM CEST
CALL HOME RETRY #0 of 12 on Oct 16, 2019 3:15:47 PM CEST.

REFCODE BE83CB93..... <=== system reference code (SRC)

SERVICEABLE EVENT TEXT

Device adapter reset reached threshold, adapter fenced. ... <=== Description of Problem

FRU group MEDIUM FRU class FRU
FRU Part Number 01LT624 FRU CCIN DAQN
FRU Serial Number 0095G725
FRU Location Code U1500.1B2.RJAAL2Y-P1-C6
FRU Previously Replaced No
FRU Previous PMH N/A

***** END OF NOTE LOG *****

For more information about planning the connections that are needed for HMC installations, see Chapter 6, “IBM DS8900F Management Console planning and setup” on page 167.

For more information about setting up SNMP notifications, see Chapter 12, “Monitoring and support” on page 423.

Remote support

Remote support provides the ability of IBM Support personnel to remotely access the DS8900F. This capability can be configured at the HMC, and access is through Assist On-site (AOS) or by IBM Remote Support Center (RSC).

For more information about remote support operations, see Chapter 12, “Monitoring and support” on page 423.

For more information about AOS, see *IBM Assist On-site for Storage Overview*, REDP-4889.

3.5 RAS on the storage system

The DS8900F was designed to safely store and retrieve large amounts of data. Redundant array of independent disks (RAID) is an industry-wide method to store data on multiple physical disks to enhance data redundancy. Many variants of RAID are used today. The DS8900F system supports RAID 6, RAID 10, and RAID 5 (by Request for Price Quotation (RPQ) only).

Note: Due to the added resiliency of RAID 6, RAID 5 is not recommended and only supported by RPQ.

3.5.1 RAID configurations

The following RAID configurations are supported on DS8900F:

- ▶ 5+P+Q+S RAID 6 configuration: The array consists of five data drives and two parity drives. The remaining drive on the array site is used as a spare.
- ▶ 6+P+Q RAID 6 configuration: The array consists of six data drives and two parity drives.
- ▶ 3+3+2S RAID 10 configuration: The array consists of three data drives that are mirrored to three copy drives. Two drives on the array site are used as spares.
- ▶ 4+4 RAID 10 configuration: The array consists of four data drives that are mirrored to four copy drives.
- ▶ 6+P+S RAID 5 configuration (by RPQ only): The array consists of six data drives and one parity drive. The remaining drive of the array site is used as a spare.
- ▶ 7+P RAID 5 configuration (by RPQ only): The array consists of seven data drives and one parity drive.

Note: The following characteristics refer to RAID:

- ▶ Spare drives are globally available to the flash RAID controller pair.
- ▶ The P and Q indicators do not mean that individual drives are dedicated to holding the parity bits for the array, but rather, they designate the equivalent capacity of the parity drives. By design, RAID 6 and RAID 5 both employ a rotating parity architecture so that no single drive is always involved in every write operation. The data and parity stripes are distributed among the member drives of the array to provide optimum write performance.

IBM Storage Modeler is an easy-to-use web tool that is available only to IBM personnel and Business Partners to help with capacity planning for physical and usable capacities that are based on installation drive capacities and quantities in intended RAID configurations.

RAID 6 is the default when creating new arrays by using the DS Storage Manager GUI.

Important: The following restrictions apply:

- ▶ A Request for Price Quotation (RPQ) / Storage Customer Opportunity Request (SCORE) is required to use RAID 5.
- ▶ Within one High-Performance Flash Enclosure (HPFE) Gen2 pair of six array sites, a RAID intermix is allowed, but no intermix of high-performance drives (Flash Tier 0) with high-capacity drives (Flash Tier 1 or Flash Tier 2) is supported.

For the latest information about supported RAID configurations and to request an RPQ / SCORE, contact your IBM SSR.

3.5.2 Drive path redundancy

Each flash drive inside an HPFE in the DS8900F is attached to two SAS switch expanders, which are contained within the Enclosure Service Modules (ESMs) in the flash drive enclosure. Figure 3-6 shows the redundancy features of the DS8900F switched SAS drive architecture.

Each flash drive has two separate connections to the enclosure backplane. This configuration allows a flash drive to be simultaneously attached to both SAS expander switches. If either ESM is removed from the enclosure, the SAS expander switch in the remaining ESM retains the ability to communicate with all the flash drives and both flash RAID controllers in the DA pair. Similarly, each DA has a path to each switch, so it can also tolerate the loss of a single path. If both paths from one DA fail, it cannot access the switches. However, the partner DA retains connectivity to all drives in the enclosure pair.

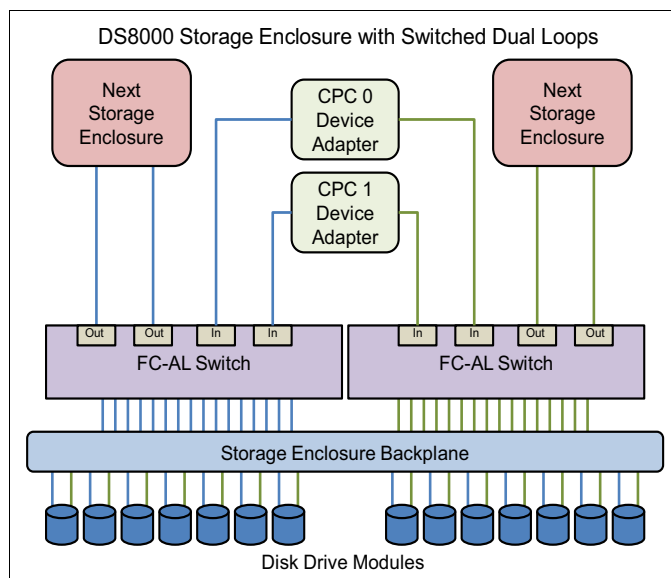


Figure 3-6 Flash Drive Enclosure Paths

For more information about the drive subsystem of the DS8900F, see 2.5, “Flash drive enclosures” on page 56.

3.5.3 Flash RAID controller redundancy

Flash RAID controllers are always installed in pairs and connect to a pair of HPFEs, which are also always installed in pairs. Each controller provides four ports that are connected across each of the four ESMs in the flash enclosures, for a total of eight paths.

The arrays are balanced between the flash enclosures to provide redundancy and performance. Both flash RAID controllers can access all arrays within the DA pair. Each controller in a DA pair is installed in different I/O enclosures, and each has allegiance to a different CPC.

Figure 3-7 on page 91 shows the connections for the DA and flash enclosure pair.




Figure 3-7 Dual paths on a flash drive enclosure

3.5.4 Predictive failure analysis

The flash drives that are used in the DS8900F incorporate PFA, and can anticipate certain forms of failures by keeping internal statistics of read/write errors. If the error rates exceed predetermined threshold values, the drive is nominated for replacement. Because the drive did not yet fail, data can be copied directly to a spare drive by using the technique that is described in 3.5.10, “Smart Rebuild” on page 95. This copy ability avoids using RAID recovery to reconstruct all the data on to the spare drive.

3.5.5 Disk scrubbing

The DS8900F periodically reads data on a flash drive. This reading is designed to occur without interfering with application performance. It is called *background data scrubbing*.

If ECC detects correctable bad bits, the bits are corrected immediately. This ability reduces the possibility of multiple bad bits accumulating in a block beyond the ability of ECC to correct them. If a block contains data that is beyond ECC’s ability to correct, RAID is used to regenerate the data and write a new copy onto a spare block or cell of the flash drive. This scrubbing process applies to flash drives that are array members and spares.

Data scrubbing can proactively relocate data, which reduces the probability of data reread impact. Data scrubbing does this relocation before errors add up to a level beyond error correction capabilities.

3.5.6 RAID support

Arrays can be configured as RAID 6, RAID 10, or RAID 5 (depending on the drive type).

Important: RAID 6 is now the default and preferred setting for the DS8900F. RAID 5 can be configured with exceptions, but it is not recommended and requires an RPQ. On high-capacity tiers (Tier 1 and Tier 2) RAID 5 is not allowed at all. RAID 10 continues to be an option for all-flash drive types.

The DS8900F uses the idea of *rotating parity*, which means that no single drive in an array is dedicated to holding parity data, which makes the drive active in every I/O operation. Instead, the drives in an array rotate between holding data stripes and holding parity stripes, balancing out the activity level of all drives in the array.

Spare drives

An HPFE Gen2 pair in a DS8900F can contain up to six array sites. Each array site contains eight flash drives, and the HPFE Gen2 pair has two spare flash drives for each enclosure pair. The first two array sites on a flash RAID controller (DA) pair have a spare that is assigned, and the rest of the array sites have no spare that is assigned if all flash drives are the same capacity. The number of required spare drives per flash enclosure pair applies to all available RAID levels.

3.5.7 RAID 6 overview

The DS8900F supports RAID 6 protection. RAID 6 presents an efficient method of data protection in double failure scenarios, such as two drive failures, two coincident media errors, or a hardware failure combined with a media error. RAID 6 protection provides greater fault tolerance than RAID 5 while consuming less raw drive capacity than RAID 10.

Note: RAID 6 is the default and preferred array configuration in DS8900F.

RAID 6 provides around a 1,000 times improvement over RAID 5 for impact risk. RAID 6 allows more fault tolerance by using a second independent distributed parity scheme (dual parity). Data is striped on a block level across a set of drives, similar to RAID 5 configurations. The second set of parity is calculated and written across all the drives, and allows reconstruction of the data even when two drives fail. The striping is shown in Figure 3-8.

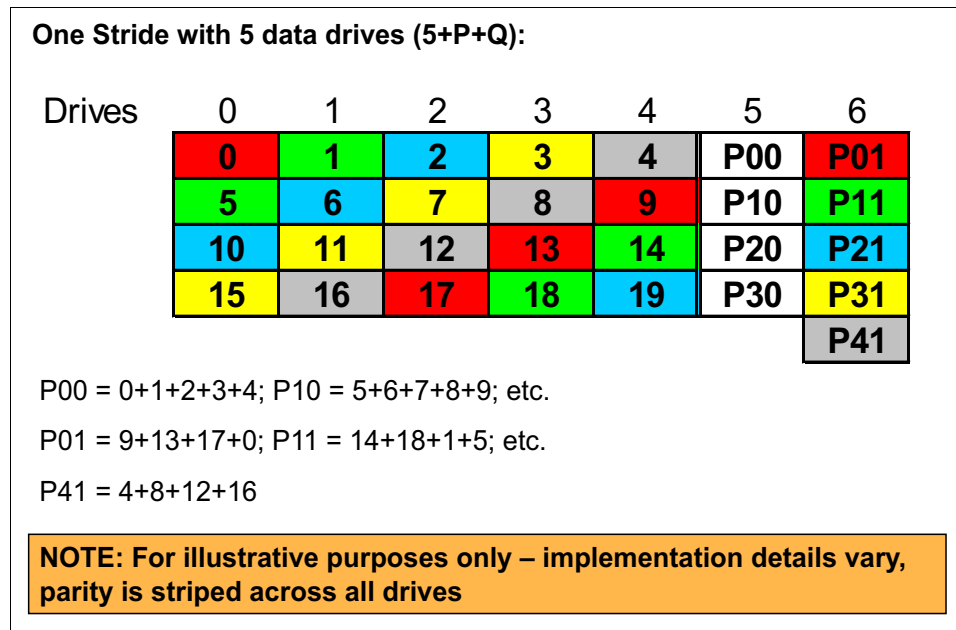


Figure 3-8 Illustrating one RAID 6 stripe on a 5+P+Q+S array

RAID 6 is best used with large-capacity drives because they have a longer rebuild time. One risk is that longer rebuild times increase the possibility that a second drive error occurs during the rebuild window. Comparing RAID 6 to RAID 5 performance gives about the same results on reads.

For random writes, the throughput of a RAID 6 array is only two-thirds of a RAID 5 due to the extra parity handling. Workload planning is important before considering RAID 6 for write-intensive applications, including CS targets. In the case of high random-write ratios, RAID 10 can be the better choice.

When RAID 6 is sized correctly for the I/O demand, it is a considerable reliability enhancement, as shown in Figure 3-8 on page 92.

RAID 6 implementation in the DS8900F

A RAID 6 array in one array site of a DS8900F can be built in one of the following configurations:

- ▶ In a seven-drive array, two drives are always used for parity, and the eighth drive of the array site is needed as a spare. This type of RAID 6 array is referred to as a 5+P+Q+S array, where P and Q stand for parity and S stands for spare.
- ▶ A RAID 6 array, which consists of eight drives, is built when all necessary spare drives are configured for the DA pair. An eight-drive RAID 6 array also always uses two drives for parity, so it is referred to as a 6+P+Q array.

Drive failure with RAID 6

When a drive fails in a RAID 6 array, the DA starts to reconstruct the data of the failing drive on to one of the available spare drives. An algorithm determines the location of the spare drive to use, depending on the capacity and the location of the failed drive. After the spare drive replaces a failed drive in a redundant array, the recalculation of the entire contents of the new drive is run by reading the corresponding data and parity in each stripe from the remaining drives in the array. This data is then written to the spare drive.

During the rebuilding of the data on the new drive, the DA can still handle I/O requests from the connected hosts to the affected array. Performance degradation might occur during the reconstruction because DAs and path resources are used to do the rebuild. Because of the dual-path architecture of the DS8900F, this effect is minimal. Additionally, any read requests for data on the failing drive require data to be read from the other drives in the array, and then the DA reconstructs the data.

Any subsequent failure during the reconstruction within the same array (second drive failure, second coincident medium errors, or a drive failure and a medium error) can be recovered without data loss.

Performance of the RAID 6 array returns to normal when the data reconstruction on the spare drive is complete. The rebuild time varies, depending on the capacity of the failed drive and the workload on the array and the DA. The completion time is comparable to a RAID 5 rebuild, but slower than rebuilding a RAID 10 array in a single drive failure.

3.5.8 RAID 10 overview

RAID 10 provides high availability (HA) by combining features of RAID 0 and RAID 1. RAID 0 optimizes performance by striping volume data across multiple drives. RAID 1 provides *drive mirroring*, which duplicates data between two drives. By combining the features of RAID 0 and RAID 1, RAID 10 provides more optimization for fault tolerance. Data is striped across half of the drives in the RAID 1 array. The same data is also striped across the other half of the array, which creates a mirror. Access to data is preserved if one drive in each mirrored pair remains available.

RAID 10 offers faster data reads and writes than RAID 6 or RAID 5 because it does not need to manage parity. However, with half of the drives in the group used for data and the other half mirroring that data, RAID 10 arrays have less usable capacity than RAID 6 or RAID 5 arrays.

RAID 10 is commonly used for workloads that require the highest performance from the drive subsystem. With RAID 6, each front-end random write I/O might theoretically lead to six back-end I/Os, including the parity updates (RAID penalty), but this number is four for RAID 5 and only two for RAID 10 (not counting cache optimizations). A typical use case for RAID 10 is for workloads with a high random-write ratio. Either member in the mirrored pair can respond to the read requests.

RAID 10 implementation in DS8900F

In the DS8900F, the RAID 10 implementation is achieved by using six or eight drives. If spares must be allocated from the array site, six drives are used to make a three-drive RAID 0 array, which is then mirrored to a three-drive array (3x3). If spares do not need to be allocated, eight drives are used to make a four-drive RAID 0 array, which is then mirrored to a four-drive array (4x4). For the required number of spares per flash drive enclosure pair, see “Spare drives” on page 92.

Drive failure with RAID 10

When a flash drive fails in a RAID 10 array, the DA rejects the failing drive and takes a hot spare into the array. Then, data is copied from the good drive to the hot spare drive. The spare that is used is chosen based on an algorithm that looks at the location of the spares and the size and location of the failed drive. Remember, a RAID 10 array is effectively a RAID 0 array that is mirrored. Therefore, when a drive fails in one of the RAID 0 arrays, you can rebuild the failed drive by reading the data from the equivalent drive in the other RAID 0 array.

While this data copy is occurring, the DA can still service read/write requests to the array from the hosts. Performance might degrade while the copy operation is in progress because DAs and path resources are used to rebuild the RAID 1 pair. Because a good drive is available, this effect is minimal. Read requests for data on the failed drive likely are not affected because they are all directed to the good copy on the mirrored drive. Write operations are not affected.

Performance of the RAID 10 array returns to normal when the data copy onto the spare drive completes. The time that is taken for rebuild can vary, depending on the capacity of the failed drive and the workload on the array and the DA.

Compared to RAID 5 or RAID 6, RAID 10 rebuild completion time is faster because rebuilding a RAID 5 or RAID 6 array requires several reads on the remaining stripe units plus one parity operation for each write. However, a RAID 10 configuration requires one read and one write (essentially, a direct copy).

3.5.9 RAID 5 implementation in DS8900F

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

Important: RAID 5 can be configured for Tier 0 flash drives of less than 1 TB, but this configuration is not recommended, and requires a risk acceptance and an RPQ for high-performance flash drives. Tier 0 flash drive sizes larger than 1 TB (not Tier 1 and Tier 2 high-capacity flash drives) can be configured by using RAID 5, but require an RPQ and an internal control switch to be enabled.

An array site with a spare creates a RAID 5 array that is 6+P+S (where the P stands for parity and S stands for spare). The other array sites on the DA pair are 7+P arrays.

3.5.10 Smart Rebuild

Smart Rebuild is a function that is designed to help reduce the possibility of secondary failures and data loss of RAID arrays. It can be used to rebuild a RAID 6 array when certain drive errors occur and a normal determination is made that it is time to use a spare to proactively replace a failing flash drive. If the suspect drive is still available for I/O, it is kept in the array rather than being rejected as under a standard RAID rebuild. A spare is brought into the array, as an extra member, concurrently.

The suspect drive and the new member-spare are set up in a temporary RAID 1 association, allowing the troubled drive to be duplicated onto the spare rather than running a full RAID reconstruction (rebuild) from data and parity. The new member-spare is then made a regular member of the array and the suspect drive is rejected from the RAID array. The array never goes through an n-1 stage in which it might suffer a complete failure if another drive in this array encounters errors. The result saves substantial time and provides a new level of availability that is not available in other RAID products.

Smart Rebuild is not applicable in all situations, so it is not always used. Smart Rebuild runs only for healthy RAID arrays. If two drives with errors are in a RAID 6 configuration, or if the drive mechanism failed to the point that it cannot accept any I/O, the standard RAID rebuild procedure is used for the RAID array. If communications across a drive fabric are compromised, such as an SAS path link error that causes the drive to be bypassed, standard RAID rebuild procedures are used because the suspect drive is not available for a one-to-one copy with a spare. If Smart Rebuild is not possible or cannot complete, a standard RAID rebuild occurs.

Drive error patterns are continuously analyzed as part of the scheduled tasks that are run by the DS8900F LIC. Drive firmware is optimized to report predictive errors to the DA. At any time, when certain drive errors (following specific criteria) reach a specified threshold, the RAS LIC component starts Smart Rebuild within the hour. This enhanced technique, when it is combined with a more frequent schedule, leads to considerably faster identification of drives showing signs of imminent failure.

A fast response in fixing drive errors is vital to avoid multiple drive failures in an array, and to thus avoid potential data loss. The possibility of having an array with diminished redundancy, such as when a RAID rebuild occurs, is reduced by shortening the time when a specific error threshold is reached until Smart Rebuild is triggered, as described in the following scenarios:

- ▶ Smart Rebuild might avoid the circumstance in which a suspected drive is rejected because the Smart Rebuild process is started before rejection. Therefore, Smart Rebuild prevents the array from going to a standard RAID rebuild, during which the array has diminished redundancy.
- ▶ Because DS8900F LIC is continuously analyzing drive errors, any drive that exceeds the error threshold is detected immediately.
- ▶ The RAS LIC component starts Smart Rebuild after the Smart Rebuild threshold criteria are met. The Smart Rebuild analysis process runs every hour.

Smart Rebuild is also used to proactively rebalance member and spare drive distribution between the paths in a DA pair. Also, if a DA pair has a mix of different capacity flash drives, a larger spare may, in some cases, be taken by a smaller drive array. Smart Rebuild corrects this situation after the failing drives are replaced, and return the larger drive to the spare pool.

DS8000 Release 9.1 code provided an enhancement of the rebuild process by avoiding the rebuild of areas that are not mapped to logical volumes.

This process is performed by running a **status** command to the drives to determine whether the parity stripe is unmapped. This process prevents unnecessary writes (P/E cycles) of zeroed data to the target drive in a rebuild, allowing faster rebuild for partially allocated RAID arrays.

IBM SSRs and remote support can manually initiate a Smart Rebuild if needed, such as when two drives in an array are logging temporary media errors.

3.5.11 Spare creation

This section describes methods of spare creation.

Flash drive enclosures

When the arrays are created on a DS8900F flash drive enclosure, the LIC determines the array sites that contain spares. The first array sites on each flash RAID controller (DA) pair that are assigned to arrays contribute one or two spares (depending on the RAID option) until the DA pair has access to at least two spares, with spares initially placed on each enclosure in the pair.

A minimum of one spare is created for each array site that is assigned to an array until the following conditions are met:

- ▶ A minimum of two spares per DA pair exist.
- ▶ A minimum of two spares for the largest capacity array site on the DA pair exist.

Spare rebalancing

The DS8900F implements a spare rebalancing technique for spare drives. When a drive fails and a hot spare is taken, it becomes a member of that array. When the failed drive is repaired, DS8900F LIC might choose to allow the hot spare to remain where it was moved. However, it can instead choose to move the spare to a more optimum position. This migration is performed to better balance the spares across the two dual flash enclosure paths to provide the optimum spare location based on drive capacity and spare availability.

It might be preferable that the drive that is in use as an array member is converted to a spare. In this case, the data on that flash drive module is moved in the background onto an existing spare by using the Smart Rebuild technique. For more information, see 3.5.5, “Disk scrubbing” on page 91 and 3.5.10, “Smart Rebuild” on page 95. This process does not fail the disk that is being moved. However, the process reduces the number of available spares in the DS8900F until the migration process is complete.

In a flash drive intermix on a DA pair, it is possible to rebuild the contents of a smaller flash drive onto a larger spare drive. When the failed origin flash drive is replaced with a new drive, the DS8900F LIC moves the data back onto the recently replaced drive.

When this process completes, the smaller flash drive rejoins the array, and the larger drive becomes a spare again.

Hot-pluggable drives

Replacing a failed flash drive does not affect the operation of the DS8900F system because the drives are fully hot-pluggable. Each drive plugs into a SAS expander switch, so no path break is associated with the removal or replacement of a drive. In addition, no potentially disruptive loop initialization process occurs.

3.6 RAS on the power subsystem

This section describes the power subsystem components of the DS8900F from a RAS standpoint.

All power and cooling components that constitute the DS8900F power subsystem are fully redundant. The key element that allows this high level of redundancy is a dual power domain configuration that is formed of iPDU pairs. Dual PSUs in all major components provide a 2N redundancy for the system.

Combined with the NVDIMMs and the BPMs, which preserve the NVS write cache, the design protects the storage system in an input power failure.

The BPMs in each of the CPCs provide power to complete the movement of write data from cache memory to non-volatile flash storage if an input power loss occurs in both power domains (as described in 3.2.4, “NVS and power outages” on page 81).

The CPCs, I/O enclosures, and flash enclosures components in the frame all feature duplicated PSUs.

In addition, the ME includes redundant PSUs that provide dual power to the ME components, such as the primary and secondary HMCs, Rack Power Control (RPC) cards, and the internal Ethernet switches.

3.6.1 Power components

This section describes the following power components:

- ▶ Intelligent Power Distribution Units
- ▶ Backup Power Modules and NVDIMMs
- ▶ Flash drive enclosure power supply units
- ▶ CPC power supply units and I/O enclosure power supply units
- ▶ Rack power control cards

Intelligent Power Distribution Units

iPDUs are used to distribute power from the AC input power cords to all areas of the system. These areas include the PSUs in flash drive enclosures, CPCs, I/O enclosures, ME, and Ethernet switches.

iPDUs are installed in pairs, one in each input power domain. An iPDU module can be replaced concurrently, as described in 2.6.3, “Power domains” on page 64.

The iPDUs are firmware upgradeable and controlled and managed by the HMCs through its Ethernet interfaces.

For more information, see 2.6.2, “Intelligent Power Distribution Units” on page 59.

iPDUs support high or low voltage three-phase, and low-voltage single-phase input power. The correct power cables must be used. For more information about power cord Feature Codes, see *IBM DS8900F Introduction and Planning Guide*, GC27-9560.

Backup Power Modules and NVDIMMs

If AC input power is lost to both power domains, and is not restored within the 20 ms *ride through* time, an emergency shutdown is initiated, and the NVDIMMs copy the data from NVS to internal flash memory.

The BPMs provide the power for this emergency copy process of the NVDIMMs. They are firmware upgradeable. The condition of the BPMs is continually monitored by the CPC FSP. The BPMs have fast charge times that ensure that an empty BPM is charged and fully operational during the IML phase of the system when the storage system powers on so that no SPoF occurs. For more information, see 2.6.5, “Backup Power Modules and NVDIMM” on page 65.

The DS8900F BPMs have a 5-year lifetime. If a BPM must be replaced, the containing CPC must be set to service mode and shut down, which invokes a failover of all operations to the other CPC. Because of the high resilience of the system, the remaining CPC keeps the whole storage facility operable and in production servicing all I/Os. As a best practice, replacement should be done in a scheduled service window to avoid reduced performance and redundancy during peak workload hours. As the BPM is monitored, sufficient warning is given to schedule the service action.

Flash drive enclosure power supply units

The flash drive enclosure PSUs provide power for the drives, and they house the cooling fans for the drive enclosure. The fans draw air from the front of the frame, through the drives, and then out through the back of the frame. The entire frame cools from front to back, complying with the data center hot aisle / cold aisle cooling strategy. Redundant fans are in each PSU, and redundant PSUs are in each drive enclosure. The drive enclosure power supply can be replaced concurrently.

Each flash drive enclosure power supply plugs into two separate iPDUs, which must each be supplied by redundant independent customer power feeds.

CPC power supply units and I/O enclosure power supply units

Each CPC and I/O enclosure has dual redundant PSUs that each receive power from a designated iPDU pair. Each I/O enclosure and each CPC has its own cooling fans.

Rack power control cards

RPC cards monitor hardware conditions and provide control paths to the I/O enclosures and system LED indicators. Two RPCs are included for redundancy, and they are housed in the ME. When one RPC is unavailable, the remaining RPC performs all necessary functions. Section 2.6.1, “Rack power control cards” on page 58 explains the architectural implementation of the RPCs.

The following RPCC features are available:

- ▶ In DS8900F, all power control functions are managed by the HMCs.
- ▶ Two different buses are used for communication between each RPC and each CPC. These buses normally perform different functions, but also can maintain communication if one of the paths is failing.
- ▶ Each RPC has two firmware images. If an RPC firmware update fails, the RPC can still boot from the other firmware image. This design also provides firmware preload capability to reduce the duration of RPC firmware updates. During a firmware update, an RPC is unavailable only for the time that is required to boot from the new firmware image.
- ▶ The RPCs monitor power to the attached I/O enclosures. They also monitor environmental components, such as power, fans, and temperature for the I/O enclosures. Environmental critical and noncritical conditions can generate EPOW events. Critical events trigger the correct signals from the hardware to the affected components to prevent any data loss without OS or firmware involvement. Non-critical environmental events are also logged and reported.

3.6.2 Line power loss

The DS8900F uses an area of server memory as NVS. This area of memory is used to hold modified data that is not yet written to the storage drives. If power is lost to both input power domains, the DS8900F uses NVDIMMs to preserve that data. For a full explanation of the NVS and cache operation, see 3.2, “CPC failover and failback” on page 78.

3.6.3 Line power fluctuation

If a power fluctuation occurs that causes a momentary interruption to power (often called a *brownout*), the DS8900F tolerates this condition for approximately 20 ms. If this period is exceeded, the flash enclosure, CPC, and other components are powered off. Independently, the NVDIMMs remain powered by the BPMs, and begin copying the contents of NVS to the flash memory to hold the data for destage when the system becomes operational. For many clients who use uninterruptible power supply (UPS) technology, brownouts are not an issue. UPS-regulated power is reliable, so more redundancy in the attached devices is often unnecessary.

3.6.4 Power control

Power control functions are performed by the HMCs, which communicate sequencing information to the service processor in each CPC, RPC, and iPDU. Power control of the DS8900F can be performed by using the Service Maintenance Console Web User Interface (WUI) or by using the DS8900F Storage Management GUI or DS CLI commands.

Figure 3-9 shows the Power Control in the Actions menu on the dashboard of the Storage Management GUI.

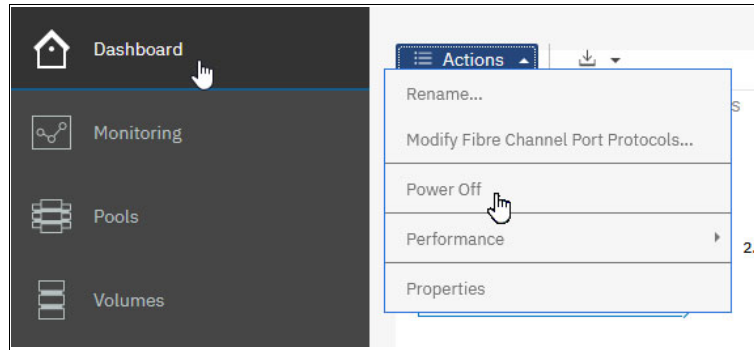


Figure 3-9 DS8900F power control in the Storage Management GUI

Figure 3-10 shows the power control settings window of the Storage Management GUI.

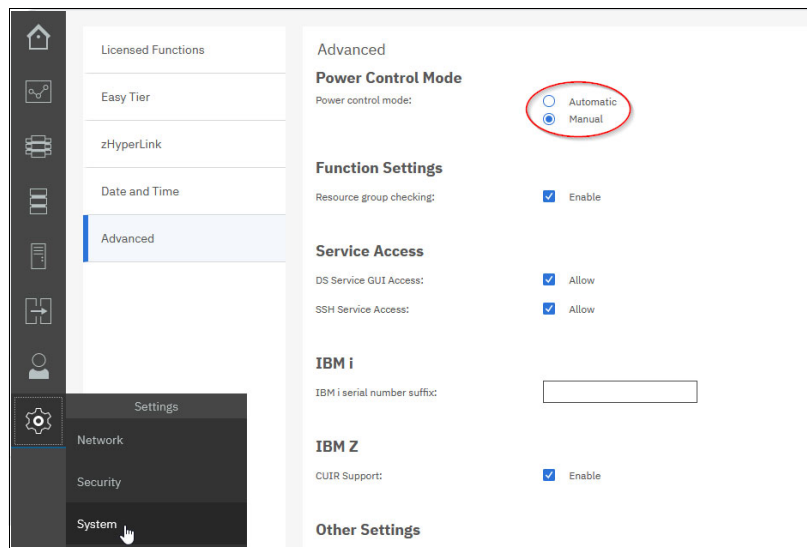


Figure 3-10 DS8900F modify power control settings from the Storage Management GUI

In addition, the following switches in the ME of a DS8900F are accessible when the ME cover is open:

- ▶ The local mode jumper connector on the local remote switch card is for service use only. There is a plug to set the system to local (ME) power control mode.
- ▶ The local power on / local force power off switch (white switch) is also on the local remote switch card in the ME. This switch can manually power on or force power off the complete system if the local remote switch card is in local power control mode. When the local / remote switch card is in remote power control mode (nothing plugged in the local mode jumper connector), the HMCs are in control of power-on / power-off (this condition is the default for client usage).
- ▶ Powering off the storage system with the white switch is a forceful shutdown. It includes the procedure of moving NVS data to the flash portion of the NVDIMMs, which must be destaged on the next system power-on and start.

Important: The local / remote power off switch (white switch) must be used only by service personnel. The switch can be used only under certain circumstances and as part of an action plan or problem determination that is performed by an IBM SSR.

3.7 Other features

Many more features of the DS8900F enhance RAS. Several of these features are described in this section.

3.7.1 Internal network

Each DS8900F base frame contains two gigabit Ethernet (GbE) switches supporting a fully redundant pair of private management networks. Each CPC in the DS8900F has a connection to each switch. Each HMC has a connection to each switch. This configuration means that if a single Ethernet switch fails, all communication from the HMCs to other components in the storage system continue to function by using the alternative private network. If a model E96 is added to the DS8980F or DS8950F, two extra Ethernet switches are added to the base frame to provide internal private network connections to the expansion frame iPDUs. Section 6.1.3, “Private and Management Ethernet networks” on page 171 explains the design of the internal network in more detail.

Note: The Ethernet switches that are used internally in DS8900F are for private network communication only. No external connection to the private networks is allowed. Client connectivity to the DS8900F is allowed only through the provided external customer HMC Ethernet connectors (eth2 and eth1) at the rear of the base frame.

3.7.2 Earthquake resistance

The Earthquake Resistance Kit is an optional seismic kit for stabilizing the storage system frame so that the frame complies with IBM earthquake resistance standards. It helps to prevent personal injury and increases the probability that the system is available following an earthquake by limiting potential damage to critical system components.

Storage system frames with this optional seismic kit include hardware at the bottom of the frame that secures it to the floor. Depending on the flooring in your environment (specifically, non-raised floors), installation of the required floor mounting hardware might be disruptive. This kit must be special-ordered for the DS8900F. The kit is not available for the rack-mountable DS8910F model 993. For more information, contact your IBM SSR.

3.7.3 IBM Certified Secure Data Overwrite

IBM Certified Secure Data Overwrite (SDO) is a process that provides a secure erasure of all data in a DS8900F storage system. Before you perform a secure data erasure, you must remove all logical configuration. Encryption groups, if configured, must also be disbanded. Then, the process is initiated by the IBM SSR. The process continues unattended until it completes. This process can take a full day to complete.

The storage system also overwrites the areas that are usually not accessible and used only internally by the disk.

As illustrated in Figure 3-11, the data becomes unintelligible after this process.

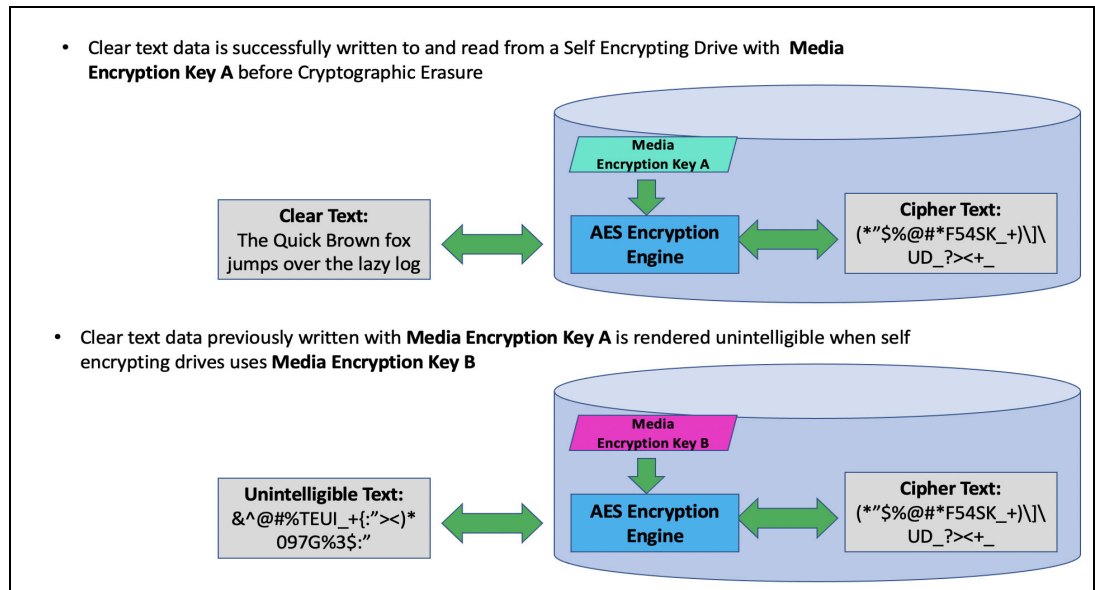


Figure 3-11 Secure Data Erasure data workflow

Flash drives secure erasure

For flash-based drives or media, a combination of crypto-erase is run by zeroing any stored key information within flash drives, resetting the flash tables in RAM, and issuing a block erase to every flash block on the drive (format operation).

CPC and HMC

CPC drives are cleared through a single-pass overwrite, which is in accordance with National Institute of Standards and Technology (NIST) SP 800-88 Rev. 1. The HMC disk drives are also overwritten with a single pass. There is no customer data on the HMC drives because this process clears all diagnostics and support trace information.

NVDIMM

The NVDIMMs are cleared by applying a single-pass overwrite in accordance with NIST SP-800-881. This process is run in parallel on both CPCs.

Process overview

The SDO process is summarized in these steps:

1. After the logical configuration is removed, SDO is started from the primary HMC.
2. The primary HMC performs a dual cluster restart.
3. The crypto-erase and format of the flash drives is started.
4. The overwrite of the CPC hard disk drives (HDDs) are started in parallel (with each other and with the above).
5. The overwrite of the secondary HMC is started in parallel.
6. Both CPCs are restarted and the NVDIMMs are cleared.
7. After the overwrite of the CPC and secondary HMC HDDs is complete, the primary HMC HDD is overwritten.
8. The certificate is generated.

Certificate

The certificate provides written verification, by drive or flash drive serial number, of the full result of the overwrite operations. You can retrieve the certificate by using DS CLI, or your IBM SSR can offload the certificate to removable media and provide it to you. Example 3-3 shows a sample SDO certificate.

Example 3-3 SDO sample certificate

Copyright (C) 2008, 2020 by International Business Machines Corporation
All Rights Reserved

Secure Data Overwrite Service for the IBM System Storage DS8900F

Certificate of Completion

IBM Corporation hereby confirms that:

1. IBM and the Customer entered into an agreement under which the Customer retained IBM to provide Secure Data Overwrite Service for the below mentioned system
2. IBM performed such Secure Data Overwrite Service as set forth herein
3. IBM provided the Customer with a report, a copy of which is attached to this certificate to identify the specific drives and the level of Secure Data Overwrite Services applied on each drive.

Machine Type, Model Serial #: 5331-996*75NHxxx
System Serial #: 75NHxxx

Customer Data Drive Overwrite Option:
Single-pass Cryptographic Erase followed by a Sanitize Block Erase

Date of Service Completion: Mar 10, 2021

In all cases, the successful completion of all erasure commands is a prerequisite for successful erasure.

Flash module (shown as 2.5" FLASH-FDE) were PURGED in accordance with NIST SP-800-88R1 for flash-based media, by issuing the sanitize command, which performs a crypto erase followed by block overwrite.

NVDIMMs NAND Flash blocks were CLEARED in accordance with NIST SP-800-88R1 for flash-based media, by applying a single overwrite pattern of 0x00. After the blocks are cleared, the data was read back to verify that the contents are erased. The overwrite and verification was performed by using vendor provided tools/methods.

CPC drives were CLEARED in accordance with NIST SP-800-88R1 for magnetic disks, by applying a single pass overwrite pattern of 0x00. Random samples, the first two sectors, and the last 10000 sectors were read back and verified to match the data written.

HMC flash base media drives were NOT securely erased. This device does not contain customer data, but the partition containing all trace data and diagnostic dumps was overwritten with single pass overwrite pattern of 0x00.

Scope

=====

This report covers the secure data overwrite service that is performed on the DS8900F storage system with the serial number 75NH430

Drive Types Table

=====

Drive Type	Drive Block Type	Drive Capacity (GB)	Number of drives installed	Number of drives available for secure overwrite process
2.5" FLASH-FDE	528	400	16	16

The Drive Types Table provides information about each drive type that is installed on the DS8000 system.

- a) Drive Type: This identifies that the drive is solid-state class of full disk encryption drive that is, 2.5" FLASH-FDE.
- b) Drive block type: This identifies that the drive block consists of 528 bytes.
- c) Drive Capacity: This identifies the specified drive type's capacity in GB.

- d) Number of drives installed: This identifies the number of drives that are installed of a given drive type and capacity.
- e) Number of drives available for secure overwrite process: This identifies the number of drives of a given drive type and capacity that is available for secure data overwrite.

Customer Data Drive Overwrite Results

=====

This section covers the devices that are used to store customer data (and associated metadata) both of which are subject to erasure.

- Disk Type - All these devices are flash memory-based and are labeled as FLASH-FDE
- Disk Serial# - Manufacturer assigned serial number visible on the device case
- WWNN - Device WWNN
- Drive Location - Rack, Enclosure, and slot where the device is installed.
- Overwrite Status - The success or failure of the overwrite operation
- Sector Defect Count - Always zero for these devices.

Customer Data Drive Table

Disk Type	Disk Serial#	WWNN	Drive Location	Overwrite Status	Sector Defect Count
2.5" FLASH-FDE	ZAJ1304E	5000C50030172544	R1-F08-D8	Successful	0
2.5" FLASH-FDE	ZAJ131SE	5000C5003017349C	R1-F08-D7	Successful	0
2.5" FLASH-FDE	ZAJ131ZE	5000C50030173AB8	R1-F08-D6	Successful	0
2.5" FLASH-FDE	ZAJ13243	5000C500301736D8	R1-F08-D5	Successful	0
2.5" FLASH-FDE	ZAJ132AT	5000C50030173A58	R1-F08-D3	Successful	0
2.5" FLASH-FDE	ZAJ132L2	5000C5003017383C	R1-F08-D4	Successful	0
2.5" FLASH-FDE	ZAJ132GW	5000C500301735B4	R1-F08-D2	Successful	0
2.5" FLASH-FDE	ZAJ132AQ	5000C50030173648	R1-F08-D1	Successful	0
2.5" FLASH-FDE	ZAJ15HE5	5000C500302283B4	R1-F07-D8	Successful	0
2.5" FLASH-FDE	ZAJ15HWY	5000C5003022817C	R1-F07-D7	Successful	0
2.5" FLASH-FDE	ZAJ15HPK	5000C500302281E0	R1-F07-D6	Successful	0
2.5" FLASH-FDE	ZAJ15KMO	5000C5003023CB00	R1-F07-D5	Successful	0
2.5" FLASH-FDE	ZAJ15HT5	5000C5003022849C	R1-F07-D2	Successful	0
2.5" FLASH-FDE	ZAJ156WJ	5000C50030228068	R1-F07-D1	Successful	0
2.5" FLASH-FDE	ZAJ15L1R	5000C5003023CB88	R1-F07-D3	Successful	0
2.5" FLASH-FDE	ZAJ15KJD	5000C5003023CAE0	R1-F07-D4	Successful	0

CPC Drive Override Results

=====

This section covers the devices on the processors that are used to store the operating system, configuration data and trace data on the Central Processor Complex (CPC) servers.

- Processor Complex # - This indicates the CPC in the DS8900F cluster
- Hdisk Number - The operating system assigned identifier
- CPC Drive Serial Number - Manufacturer assigned serial number visible on the device case
- Overwrite Status - The success or failure of the overwrite operation
- Completion Date - Completion timestamp

CPC Drive Table

Processor Complex #	hdisk Number	CPC Drive Serial Number	Overwrite Status	Completion Date
CPC 0	hdisk0	WAE1045Q	Successful	2021/03/09 19:11:46
CPC 0	hdisk1	WAE10N39	Successful	2021/03/09 20:10:31
CPC 1	hdisk0	0TJ55JLP	Successful	2021/03/09 19:13:16
CPC 1	hdisk1	WAE104DZ	Successful	2021/03/09 20:12:32

Hardware Management Console (HMC) Drives

=====

This section covers the devices on the processors that are used to store the operating system,

configuration data and trace data on the Hardware Management Console (HMC).

As noted above, these devices were NOT erased and only the partition containing logs and dumps were deleted.

- HMC Type - Indicates whether this is the first or optional second HMC
- HMC Drive Serial Number - Manufacturer assigned serial number visible on the device case
- Overwrite Status - The success or failure of the overwrite operation
- Completion Date - Completion timestamp
- HMC Drive Type - Always SSD for these systems

Hardware Management Console (HMC) Override Results

```
=====
```

HMC Type	HMC Drive Serial Number	SDO Results	Completion Date	HMC Drive Type Hard Disk Drive/ SSD
First Management Console	N/A	Successful	03/10/21-06:33:51	SSD
Secondary Management Console	N/A	Successful	03/10/21-03:49:39	SSD

```
=====
```

Non-Volatile Dual In-Line Memory Module (NVDIMM)

```
=====
```

This section covers the devices that are used to store customer data when system goes through emergency power off and the device is subject to erasure.

- NVDIMM Location Code - Central Processor Complex (CPC) and slot where the device is installed
- Serial Number- Manufacturer assigned serial number visible on the device
- NVDIMM Capacity- Device capacity in GB
- Overwrite Status- The success or failure of the overwrite operation
- Completion Date - Completion timestamp

NVDIMM Secure Erase Results

```
=====
```

NVDIMM Location Code	Serial Number	NVDIMM Capacity	Overwrite Status	Completion Date
U78D2.001.WZS03ED-P1-C22	YH10DP98G0E7	32 GB	Successful	03/09/21 19:42:56
U78D2.001.WZS03ED-P1-C36	YH10DP98G0EX	32 GB	Successful	03/09/21 19:42:56
U78D2.001.WZS03RG-P1-C22	YH30DP944055	32 GB	Successful	03/09/21 19:42:58
U78D2.001.WZS03RG-P1-C36	YH30DP94400R	32 GB	Successful	03/09/21 19:42:58

```
=====
```




Virtualization concepts

This chapter describes virtualization concepts as they apply to the IBM DS8900F.

This chapter covers the following topics:

- ▶ Virtualization definition
- ▶ Benefits of virtualization
- ▶ Abstraction layers for drive virtualization
- ▶ Extent pools

4.1 Virtualization definition

For the purposes of this chapter, *virtualization* is the abstraction of a physical drive to one or more logical volumes. This *virtual drive* is presented to hosts and systems as though it is a physical drive.

4.2 Benefits of virtualization

Virtualization in IBM DS8900F includes the following benefits:

- ▶ Flexible logical volume configuration:
 - Multiple redundant array of independent disks (RAID) types (RAID 6, RAID 5, and RAID 10). RAID 6 is preferred, and is the default for all drive types.
 - Storage types (Count Key Data (CKD) and Fixed-Block (FB) architecture) that are aggregated into separate extent pools.
 - Volumes are provisioned (allocated) from extents of the extent pool.
 - Storage pool striping.
 - Dynamically add and remove volumes.
 - Logical volume configuration states.
 - Dynamic Volume Expansion (DVE).
 - Extent space efficient (ESE) volumes for thin provisioning of FB and CKD volumes.
 - Support for small and large extents.
 - Extended address volumes (EAVs) (CKD).
 - Parallel access volumes (PAVs) across logical control units (LCUs) (SuperPAV for CKD).
 - Dynamic extent pool merging for IBM Easy Tier.
 - Dynamic Volume Relocation (DVR) for Easy Tier.
 - Easy Tier Heat Map Transfer (HMT).
- ▶ Flexible logical volume sizes:
 - CKD: Up to 1 TB (1,182,006 cylinders) by using EAVs.
 - FB: Up to 16 TB either with small or large extents (limit of 4 TB when used with Copy Services (CS)).
- ▶ Flexible number of logical volumes:
 - Up to 65280 (CKD).
 - Up to 65280 (FB).
 - 65280 total for mixed CKD + FB.
- ▶ No strict relationship between RAID ranks and logical subsystems (LSSs).
- ▶ LSS definition allows flexible configuration of the number and size of devices per LSS:
 - With larger devices, fewer LSSs can be used.
 - Volumes for a particular application can be kept in a single LSS.
 - Smaller LSSs can be defined, if required (for applications that require less storage).
 - Test systems can have their own LSSs with fewer volumes than production systems.

4.3 Abstraction layers for drive virtualization

Virtualization in the DS8900F refers to the process of preparing physical drives for storing data on logical volumes for use by attached hosts. Logical volumes are seen by the hosts as though they were physical storage. In open systems, this process is known as creating logical unit numbers (LUNs). In IBM Z, it refers to the creation of 3390 volumes.

The basis for virtualization begins with the *physical drives*, which are mounted in storage enclosures and connected to the internal storage servers. DS8900F uses only the High-Performance Flash Enclosure (HPFE) Gen2 storage enclosures. To learn more about the drive options and their connectivity to the internal *storage servers*, see 2.5, “Flash drive enclosures” on page 56.

Virtualization builds upon the physical drives as a series of layers:

- ▶ Array sites
- ▶ Arrays
- ▶ Ranks
- ▶ Extent pools
- ▶ Logical volumes
- ▶ LSSs or LCUs

4.3.1 Array sites

An *array site* is formed from a group of eight identical drives with the same capacity and drive class. The specific drives that are assigned to an array site are automatically chosen by the system at installation time to balance the array site capacity across both drive enclosures in a pair and across the connections to both storage servers. The array site also determines the drives that are required to be reserved as spares. No predetermined storage server affinity exists for array sites. Array sites are the building blocks that are used to define *arrays*.

4.3.2 Arrays

An *array* is created from one *array site*. When an array is created, its RAID level, array type, and array configuration are defined. This process is also called *defining* an array. In all IBM DS8000 series implementations, one array is always defined as using one array site.

The following RAID levels are supported in DS8900F:

- ▶ RAID 6 (default)
- ▶ RAID 10
- ▶ RAID 5 (possible, but not recommended, for flash drive sizes below 1 TB)

Each HPFE Gen2 pair can contain up to six array sites. The first set of 16 flash drives creates two 8-drive array sites. RAID 6 arrays are created by default on each array site. RAID 10 is optional for all flash drive sizes. RAID 5 is optional for flash drives smaller than 1 TB, but is not recommended, and requires risk acceptance.

A Request for Price Quotation (RPQ) is required to use of RAID 5 on flash drives greater than 1 TB (RPQ is not available for drive sizes of 4 TB and greater).

During logical configuration, RAID 6 arrays and the required number of spares are created. Each HPFE Gen2 pair has two global spares that are created from the first increment of 16 flash drives. The first two arrays to be created from these array sites are 5+P+Q. Subsequent RAID 6 arrays in the same HPFE Gen2 Pair are 6+P+Q.

Important: Using RAID 6 is recommended, and it is the default in the DS GUI. As with large drives in particular, the RAID rebuild times (after one drive failure) get ever larger. Using RAID 6 reduces the danger of data loss due to a double-drive failure. For more information, see 3.5.1, “RAID configurations” on page 89.

For more information about the sparing algorithm, see 3.5.11, “Spare creation” on page 96.

Figure 4-1 shows the creation of a RAID 6 array with one spare, which is also called a $5+P+Q+S$ array. It has a capacity of five drives for data, two drives for double distributed parity, and a spare drive. According to the RAID 6 rules, parities are distributed across all seven drives in this example.

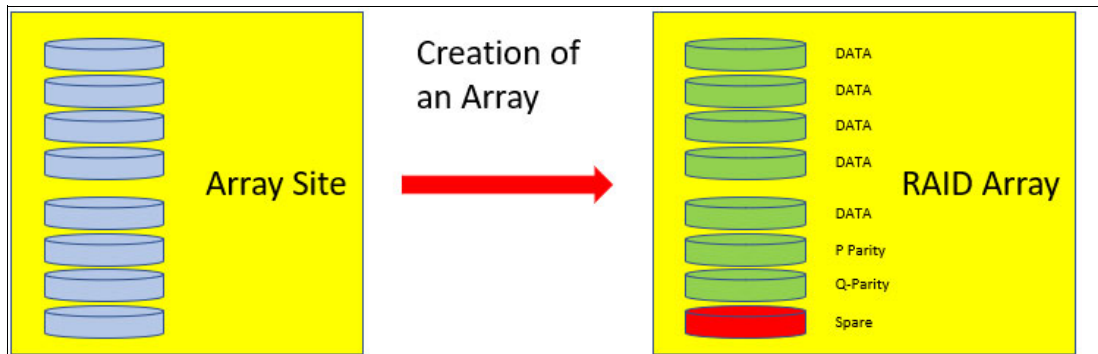


Figure 4-1 Creation of a $5+P+Q+S$ array

Depending on the selected RAID level and sparing requirements, six types of arrays are possible, as shown in Figure 4-2 on page 111.

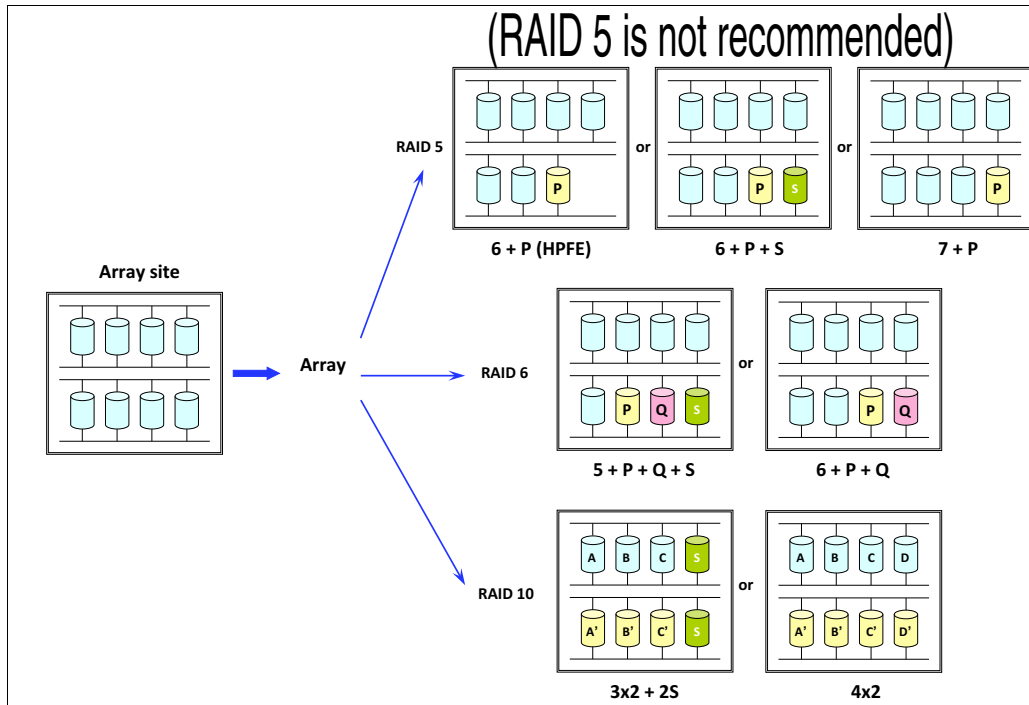


Figure 4-2 RAID array types

Tip: Larger drives have a longer rebuild time. Only RAID 6 can recover from a double disk error during a rebuild, by using the additional parity data. RAID 6 is the best choice for systems that require high availability (HA), and is the default in DS8900F.

4.3.3 Ranks

After the arrays are created, the next task is to define a rank. A *rank* is a logical representation of the physical array that is formatted for use as FB or CKD storage types. In the DS8900F, ranks are defined in a one-to-one relationship to arrays. Before you define any ranks, you must decide whether you plan to encrypt the data or not.

Encryption group

All drives that are offered in the DS8900F are Full Disk Encryption (FDE)-capable to secure all logical volume data at rest. In the DS8900, the Encryption Authorization license is included in the Base Function (BF) license group.

If you plan to use encryption for data at rest, you must define an *encryption group* before any ranks are created. The DS8900F supports only one encryption group. All ranks must be in this encryption group. The encryption group is an attribute of a rank. Therefore, your choice is to encrypt everything or nothing. If you want to enable encryption later (create an encryption group), all logical configuration must be deleted and re-created, and volume data restored.

For more information, see *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security (DS8000 Release 9.2)*, REDP-4500.

Defining ranks

When a new rank is defined, its name is chosen by the DS GUI or Data Storage Command-line Interface (DS CLI), for example, R1, R2, or R3. The rank is then associated with an array.

Important: In all DS8000 series implementations, a rank is defined as using only one array. Therefore, rank and array can be treated as synonyms.

The process of defining a rank accomplishes the following objectives:

- ▶ The array is formatted for FB data for open systems or CKD for IBM Z data. This formatting determines the size of the set of data that is contained on one drive within a stripe on the array.
- ▶ The capacity of the array is subdivided into partitions, which are called *extents*. The extent size depends on the *extent type*, whether FB or CKD. The extents are the building blocks of the logical volumes. An extent is striped across all drives in an array, and it is represented by the small squares in Figure 4-3 on page 114.
- ▶ You can choose between large extents and small extents.

An FB rank features an extent size of either 1 GB (more precisely a *gibibyte* (GiB), which is a binary gigabyte that is equal to 2^{30} bytes), called *large extents*, or an extent size of 16 mebibytes (MiB), called *small extents*.

IBM Z users or administrators typically do not deal with gigabytes or gibibytes. Instead, storage is defined in terms of the original 3390 volume sizes. A 3390 Model 3 is three times the size of a Model 1. A Model 1 features 1113 cylinders, which are about 0.946 GB.

A 3390 Model 1 (1113 cylinders) is the large extent size for CKD ranks. The CKD small extent size is 21 cylinders, which corresponds to the z/OS allocation unit for EAV volumes larger than 65520 cylinders. z/OS changes the addressing modes and allocates storage in 21 cylinder units.

An extent can be assigned to only one volume. Although you can define a CKD volume with a capacity that is an integral multiple of one cylinder or an FB LUN with a capacity that is an integral multiple of 128 logical blocks (64 KB), if you define a volume this way, you might waste the unused capacity in the last extent that is assigned to the volume.

For example, the DS8900F theoretically supports a minimum CKD volume size of one cylinder, but the volume still claims one full extent of 1113 cylinders if large extents are used or 21 cylinder for small extents. So, 1112 cylinders are wasted if large extents are used.

Note: In the DS8900F firmware, all volumes have a common metadata structure. All volumes have the metadata structure of ESE volumes, whether the volumes are thin-provisioned or fully provisioned. ESE is described in 4.4.4, “Volume allocation and metadata” on page 125.

Small or large extents

Whether you use small or large extents depends on the goals that you want to achieve. Small extents provide better capacity utilization, particularly for thin-provisioned storage. With thin-provisioned volumes, using small extents is preferred. However, managing many small extents causes some small performance degradation during initial allocation. For example, a format write of 1 GB requires one storage allocation with large extents, but 64 storage allocations with small extents. Otherwise, host performance should not be adversely affected.

4.4 Extent pools

An *extent pool* is a logical construct to aggregate the extents from a set of ranks, and it forms a domain for extent allocation to a logical volume. Originally, extent pools were used to separate drives with different revolutions per minute (RPM) and capacity in different pools that have homogeneous characteristics. You still might want to use extent pools to separate Tier 0, Tier 1, and Tier 2 flash drives, but be aware that Easy Tier does not manage data placement across extent pools.

No rank or array affinity to an internal server (central processor complex (CPC) is predefined. The affinity of the rank (and its associated array) to a server is determined when it is assigned to an extent pool. One or more ranks with the same extent type (FB or CKD) can be assigned to an extent pool.

Important: Because a rank is formatted to have small or large extents, the first rank that is assigned to an extent pool determines whether the extent pool is a pool of all small or all large extents. You *cannot* have a pool with a mixture of small and large extents. You cannot change the extent size of an extent pool.

If you want Easy Tier to automatically optimize rank utilization, configure more than one rank in an extent pool. A rank can be assigned to only one extent pool. As many extent pools as ranks can exist, but for most systems, a single pair of extent pools for each rank type (FB or CKD) provides the best overall performance.

Heterogeneous extent pools, with a mixture of Tier 0, Tier 1, and Tier 2 flash drives can take advantage of the capabilities of Easy Tier to optimize I/O throughput. Easy Tier moves data across different storage tiering levels to optimize the placement of the data within the extent pool.

With *storage pool striping*, you can create logical volumes that are striped across multiple ranks to enhance performance. To benefit from storage pool striping, more than one rank in an extent pool is required.

Storage pool striping can enhance performance significantly. However, in the unlikely event that a whole RAID array fails, the loss of the associated rank affects the entire extent pool because data is striped across all ranks in the pool. For data protection, consider mirroring your data to another DS8000 family storage system.

When an extent pool is defined, it must be assigned the following attributes:

- ▶ Internal storage server affinity
- ▶ Extent type (FB or CKD)
- ▶ Encryption group

As with ranks, extent pools are also assigned to encryption group 0 or 1, where group 0 is non-encrypted, and group 1 is encrypted. The DS8900F supports only one encryption group, and all extent pools must use the same encryption setting that is used for the ranks.

A minimum of two extent pools must be configured to balance the capacity and workload between the two servers. One extent pool is assigned to internal server 0. The other extent pool is assigned to internal server 1. In a system with both FB and CKD volumes, four extent pools provide one FB pool for each server and one CKD pool for each server.

If you plan on using small extents for ESE volumes while retaining large extents for other volumes, you must create more pools with small extents. Small and large extents cannot be in the same pool.

Figure 4-3 shows an example of a mixed environment that features CKD and FB extent pools.

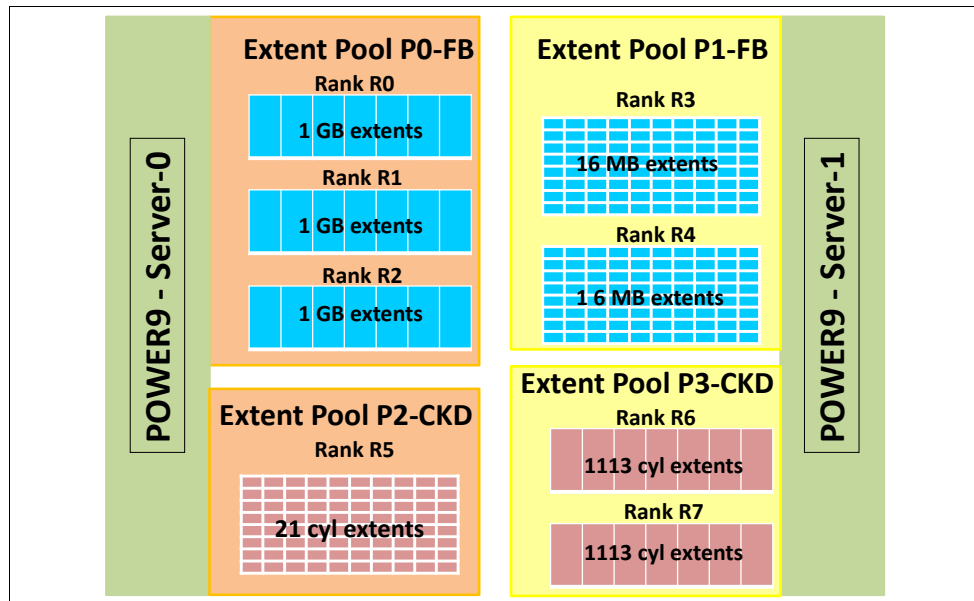


Figure 4-3 Extent pools

Extent pools can be expanded by adding more ranks to the pool. All ranks that belong to extent pools with the same internal server affinity are called a *rank group*. Ranks are organized in two rank groups. Rank group 0 is controlled by server 0, and rank group 1 is controlled by server 1.

4.4.1 Dynamic extent pool merge

Dynamic extent pool merge is a capability that is provided by the Easy Tier manual mode facility, which is shown in Figure 4-4 on page 115. It allows one extent pool to be merged into another extent pool if they meet these criteria:

- ▶ Have extents of the same size
- ▶ Have the same storage type (FB or CKD)
- ▶ Have the same DS8900F internal server affinity

The logical volumes in both extent pools remain accessible to the host systems. Dynamic extent pool merge can be used for the following reasons:

- ▶ Consolidation of two smaller extent pools with the equivalent storage type (FB or CKD) and extent size into a larger extent pool. Creating a larger extent pool allows logical volumes to be distributed over a greater number of ranks, which improves overall performance in the presence of skewed workloads. Newly created volumes in the merged extent pool allocate capacity as specified by the selected extent allocation algorithm. Logical volumes that existed in either the source or the target extent pool can be redistributed over the set of ranks in the merged extent pool by using the Migrate Volume function.
- ▶ Consolidating extent pools with different storage tiers to create a merged extent pool with a mix of storage drive technologies. This type of an extent pool is called a *multitiered pool*, and it is a prerequisite for using the Easy Tier automatic mode feature.

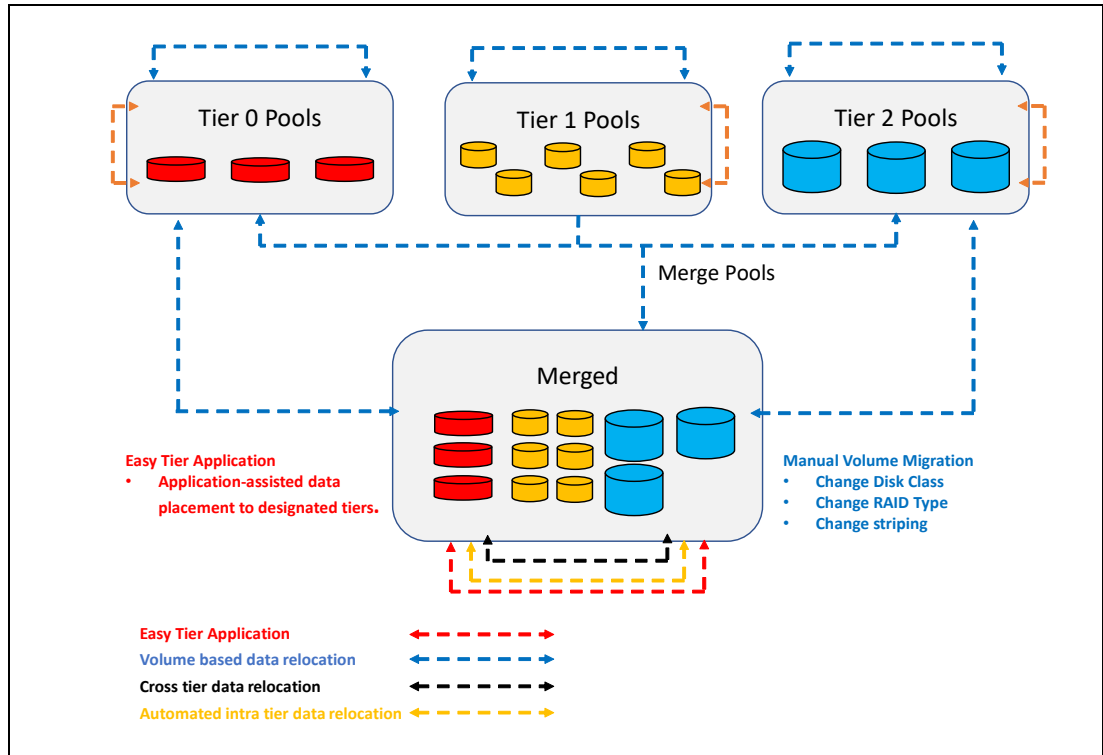


Figure 4-4 Easy Tier migration types

Important: Volume migration (or DVR) within the same extent pool is not supported in multitiered pools. Easy Tier automatic mode rebalances the volumes' extents within the multitiered extent pool automatically based on I/O activity. However, you can also use the Easy Tier application to manually place entire volumes in designated tiers. For more information, see *DS8870 Easy Tier Application*, REDP-5014.

Dynamic extent pool merge is allowed only among extent pools with the same internal server affinity or rank group. Additionally, the dynamic extent pool merge is not allowed in the following circumstances:

- ▶ If source and target pools have different storage types (FB and CKD).
- ▶ If source and target pools have different extent sizes.
- ▶ If you selected an extent pool that contains volumes that are being moved.
- ▶ If the combined extent pools include 2 PB or more of ESE effective (virtual) capacity.

For more information about Easy Tier, see *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667.

4.4.2 Logical volumes

A *logical volume* consists of a set of extents from one extent pool. The DS8900F supports up to 65,280 logical volumes (64 K CKD or 64 K FB volumes, or a mixture of both, up to a maximum of 64 K total volumes). The abbreviation 64 K is used in this section, even though it is 65,536 minus 256, which is not quite 64 K in binary.

Fixed-Block LUNs

A logical volume that is composed of FB extents is called a *LUN*. An FB LUN is composed of one or more 1 GiB (2^{30} bytes) large extents or one or more 16 MiB small extents from one FB extent pool. A LUN cannot span multiple extent pools, but a LUN can have extents from multiple ranks within the same extent pool. You can construct LUNs up to 16 TiB (16×2^{40} bytes, or 2^{44} bytes) when using large extents.

Important: DS8000 CS does not support FB logical volumes larger than 4 TiB. Do not create a LUN that is larger than 4 TiB if you want to use CS for the LUN unless the LUN is integrated as managed disks (MDisks) in an IBM SAN Volume Controller (SVC), and the LUN is using IBM Spectrum® Virtualize CS.

LUNs can be provisioned (allocated) in binary GiB (2^{30} bytes), decimal GB (10^9 bytes), or 512 or 520-byte blocks. However, the usable (physical) capacity that is provisioned (allocated) is a multiple of 1 GiB. For small extents, it is a multiple of 16 MiB. Therefore, it is a good idea to use LUN sizes that are a multiple of a gibibyte or a multiple of 16 MiB. If you define a LUN with a size that is not a multiple of 1 GiB (for example, 25.5 GiB), the LUN size is 25.5 GiB. However, 26 GiB are physically provisioned (allocated), of which 0.5 GiB of the physical storage is unusable. When you want to specify a LUN size that is not a multiple of 1 GiB, specify the number of blocks. A 16 MiB extent has 32768 blocks.

The allocation process for FB volumes is illustrated in Figure 4-5.

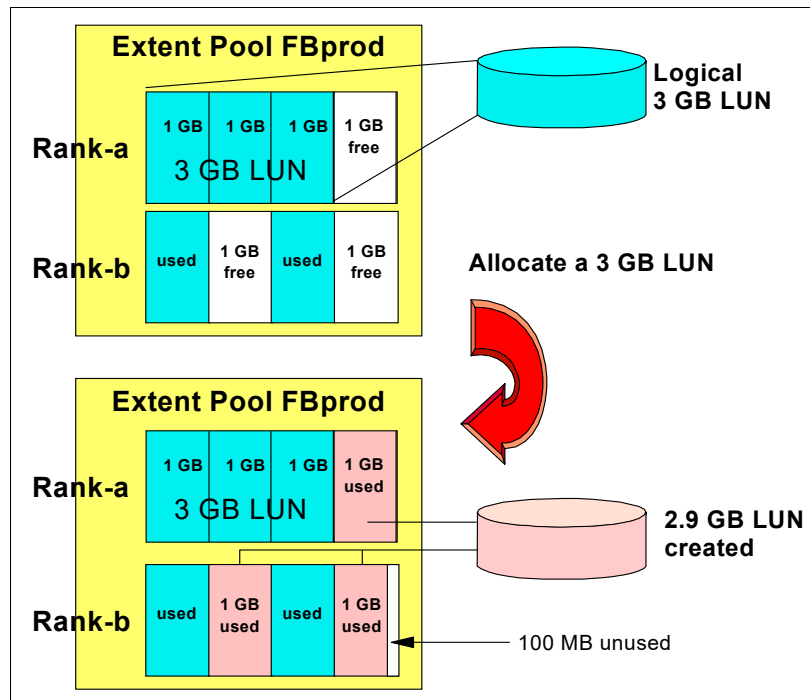


Figure 4-5 Creating an FB LUN

With small extents, wasted storage is not an issue.

An FB LUN must be managed by an LSS. One LSS can manage up to 256 LUNs. The LSSs are created and managed by the DS8900, as required. A total of 255 LSSs can be created in the DS8900.

IBM i logical unit numbers

IBM i LUNs are composed of FB 1 GiB (or 16 MiB) extents. However, special aspects must be considered with IBM i LUNs. LUNs that are created on a DS8900F are always RAID-protected. LUNs are based on RAID 10, RAID 6, or RAID 5 arrays. However, you might want to make it seem to IBM i that the LUN is not RAID-protected. This setup causes IBM i to conduct its own mirroring.

IBM i LUNs may use the **unprotected** attribute, in which case the DS8900F reports that the LUN is not RAID-protected. Selecting either the **protected** or **unprotected** attribute does not affect the RAID protection that is used by the DS8900F on the open volume.

IBM i LUNs display a 520-byte block to the host. The operating system (OS) uses eight of these bytes, so the usable space is still 512 bytes like other Small Computer System Interface (SCSI) LUNs. The capacities that are quoted for the IBM i LUNs are in terms of the 512-byte block capacity, and they are expressed in GB (10⁹). Convert these capacities to GiB (2³⁰) when you consider the effective usage of extents that are 1 GiB (2³⁰).

Important: The DS8900F supports IBM i variable volume (LUN) sizes in addition to fixed volume sizes.

IBM i volume enhancement adds flexibility for volume sizes and can optimize the DS8900F capacity usage for IBM i environments.

The DS8900F supports IBM i variable volume data types A50, which is an unprotected variable size volume, and A99, which is a protected variable size volume. For more information, see Table 4-1. DS8000 Release 9 introduced dynamic expansion of IBM i variable volumes.

Table 4-1 IBM i variable volume sizes

Model type		IBM i device size (GB)	Number of logical block addresses (LBAs)	Extents	Unusable space (GiB ¹)	Usable space%
Unprotected	Protected					
2107-050	2107-099	Variable			0.00	Variable

Example 4-1 demonstrates the creation of both a protected and an unprotected IBM i variable size volume by using the DS CLI.

Example 4-1 Creating the IBM i variable size for unprotected and protected volumes

```
dscli> mkfbvol -os400 050 -extpool P4 -name itso_iVarUnProt1 -cap 10 5413
CMUC00025I mkfbvol: FB volume 5413 successfully created.
```

```
dscli> mkfbvol -os400 099 -extpool P4 -name itso_iVarProt1 -cap 10 5417
CMUC00025I mkfbvol: FB volume 5417 successfully created.
```

When you plan new capacity for an existing IBM i system, the larger the LUN, the more data it might have, which causes more I/O operations per second (IOPS) to be driven to it. Therefore, mixing different drive sizes within the same system might lead to hot spots.

Note: IBM i fixed volume sizes continue to be supported in current DS8000 code levels. Consider the best option for your environment between fixed and variable-size volumes.

T10 Data Integrity Field support

The American National Standards Institute (ANSI) T10 standard provides a way to check the integrity of data that is read and written from the application or the host bus adapter (HBA) to the drive and back through the storage area network (SAN) fabric. This check is implemented through the Data Integrity Field (DIF) that is defined in the T10 standard. This support adds protection information that consists of a cyclic redundancy check (CRC), LBA, and host application tags to each sector of FB data on a logical volume.

A T10 DIF-capable LUN uses 520-byte sectors instead of the common 512-byte sector size. Eight bytes are added to the standard 512-byte data field. The 8-byte DIF consists of 2 bytes of CRC data, a 4-byte Reference Tag (to protect against misdirected writes), and a 2-byte Application Tag for applications that might use it.

On a write, the DIF is generated by the HBA, which is based on the block data and LBA. The DIF field is added to the end of the data block, and the data is sent through the fabric to the storage target. The storage system validates the CRC and Reference Tag and, if correct, stores the data block and DIF on the physical media. If the CRC does not match the data, the data was corrupted during the write. The write operation is returned to the host with a write error code. The host records the error and retransmits the data to the target. In this way, data corruption is detected immediately on a write, and the corrupted data is never committed to the physical media.

On a read, the DIF is returned with the data block to the host, which validates the CRC and Reference Tags. This validation adds a small amount of latency for each I/O, but it might affect overall response time on smaller block transactions (less than 4 KB I/Os).

The DS8900F supports the T10 DIF standard for FB volumes that are accessed by the Fibre Channel Protocol (FCP) channels that are used by Linux on IBM Z, or AIX. You can define LUNs with an option to instruct the DS8900F to use the CRC-16 T10 DIF algorithm to store the data.

You can also create T10 DIF-capable LUNs for OSs that do not yet support this feature (except for IBM i). Active protection is available for Linux on IBM Z, and AIX on IBM Power servers. For other distributed OSs, check their documentation.

When you create an FB LUN by running the `mkfbvo1` DS CLI command, add the option `-t10dif`. If you query a LUN with the `showfbvo1` command, the data type is FB 512T instead of the standard FB 512 type.

Important: Because the DS8900F internally always uses 520-byte sectors (to support IBM i volumes), no extra capacity is considered when standard or T10 DIF-capable volumes are used.

Target LUN: When FlashCopy for a T10 DIF LUN is used, the target LUN must also be a T10 DIF-type LUN. This restriction does not apply to mirroring.

Count Key Data volumes

An IBM Z CKD volume is composed of one or more extents from one CKD extent pool. CKD extents are of the size of 3390 Model 1, which features 1113 cylinders for large extents or 21 cylinders for small extents. When you define an IBM Z CKD volume, specify the size of the volume as a multiple of 3390 Model 1 volumes or the number of cylinders that you want for the volume.

Before a CKD volume can be created, an LCU must be defined that provides up to 256 possible addresses that can be used for CKD volumes. Up to 255 LCUs can be defined. For more information about LCUs, see 4.4.5, “Logical subsystems” on page 130.

On a DS8900F, you can define CKD volumes with up to 1,182,006 cylinders, or about 1 TB. This volume capacity is called an EAV, and it is supported by the 3390 Model A.

A CKD volume cannot span multiple extent pools, but a volume can have extents from different ranks in the same extent pool. You also can stripe a volume across the ranks. For more information, see “Storage pool striping: Extent rotation” on page 121.

Figure 4-6 shows an example of how a logical volume is provisioned (allocated) with a CKD volume.

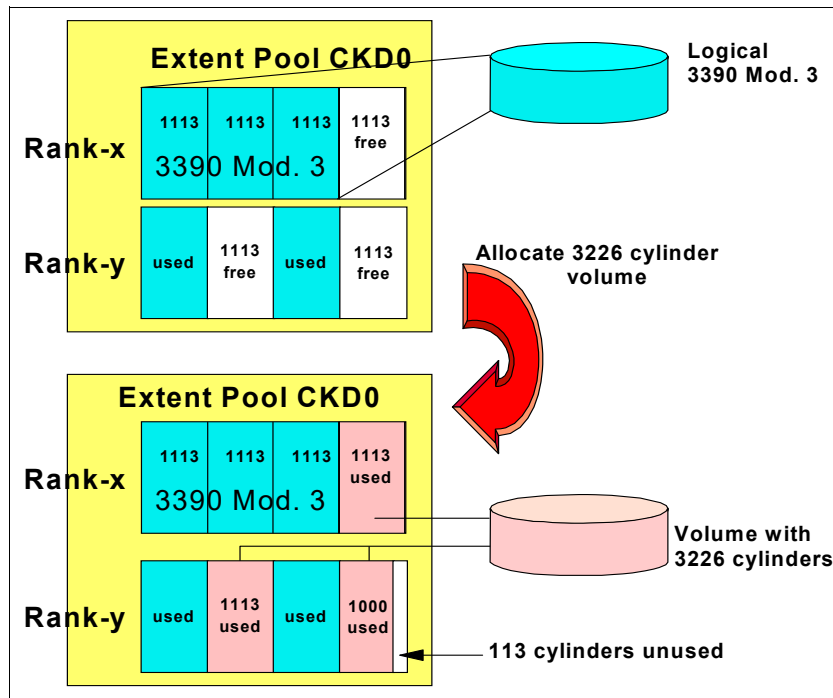


Figure 4-6 Allocating a Count Key Data logical volume

CKD alias volumes

Another type of CKD volume is the PAV *alias volume*. PAVs are used by z/OS to send parallel I/Os to the same base CKD volume. Alias volumes are defined within the same LCU as their corresponding base volumes. Although they have no size, each alias volume needs an address, which is linked to a base volume. The total of base and alias addresses cannot exceed the maximum of 256 for an LCU.

HyperPAV and SuperPAV

With HyperPAV, the alias can be assigned to access a base device within the same LCU on an I/O basis. However, the system can exhaust the alias resources that are needed within an LCU and cause an increase in I/O queue time (the queue time of software that is waiting for a device to accept the I/O). In this case, you need more aliases to access hot volumes.

Classically, to start an I/O to a base volume, z/OS can select any alias address *only* from the same LCU as the base address to perform the I/O. With SuperPAV, the OS can use alias addresses from *other LCUs* to perform an I/O for a base address.

The restriction is that the LCU of the alias address belongs to the same DS8000 server. In other words, if the base address is from an even / odd LCU, the alias address that z/OS can select must also be from an even / odd LCU. In addition, the LCU of the base volume and the LCU of the alias volume must be in the same path group. z/OS prefers alias addresses from the same LCU as the base address, but if no alias address is free, z/OS looks for free alias addresses in LCUs of the same *Alias Management Group*.

An Alias Management Group is all the LCUs that have affinity to the same DS8000 internal server and have the same paths to the DS8900F. SMF can provide reports at the Alias Management Group level.

Figure 4-7 shows how SuperPAV works.

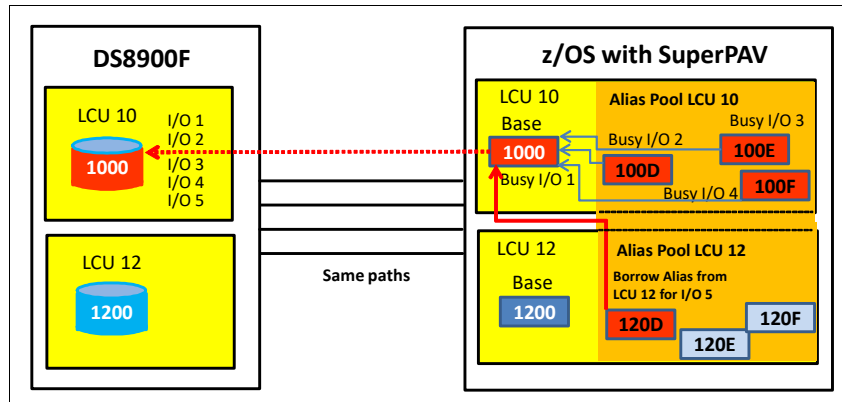


Figure 4-7 SuperPAV

Initially, each alias address must be assigned to a base address. Therefore, it is not possible to define an LCU with only alias addresses.

As with PAV and HyperPAV, SuperPAV must be enabled. SuperPAV is enabled by the `HYPERPAV=XPAV` statement in the `IECIO$xx` parmlib member or by the `SETIOS HYPERPAV=XPAV` command. The `D M=DEV(address)` and the `D M=CU(address)` display commands show whether XPAV is enabled or not. With the `D M=CU(address)` command, you can check whether aliases from other LCUs are being used (Example 4-2).

Example 4-2 Displaying information about SuperPAV (XPAV)

```

D M=CU(4E02)
IEE174I 11.26.55 DISPLAY M 511
CONTROL UNIT 4E02
CHP                65   5D   34   5E
ENTRY LINK ADDRESS 98   ..  434B ..
DEST LINK ADDRESS  FA   0D  200F 0D
CHP PHYSICALLY ONLINE Y   Y   Y   Y
  
```

```

PATH VALIDATED      Y   Y   Y   Y
MANAGED             N   N   N   N
ZHPF - CHPID       Y   Y   Y   Y
ZHPF - CU INTERFACE Y   Y   Y   Y
MAXIMUM MANAGED CHPID(S) ALLOWED = 0
DESTINATION CU LOGICAL ADDRESS = 56
CU ND               = 002107.981.IBM.75.0000000FXF41.0330
CU NED              = 002107.981.IBM.75.0000000FXF41.5600
TOKEN NED           = 002107.900.IBM.75.0000000FXF41.5600
FUNCTIONS ENABLED = ZHPF, XPAV
XPAV CU PEERS    = 4802, 4A02, 4C02, 4E02
DEFINED DEVICES
  04E00-04E07
DEFINED PAV ALIASES
  14E40-14E47

```

With cross-LCU HyperPAV, which is called SuperPAV, the number of alias addresses can further be reduced while the pool of available alias addresses to handle I/O bursts to volumes is increased.

4.4.3 Allocating, deleting, and modifying LUNs and CKD volumes

Extents of 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. 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 DS8900F allows flexibility in the management of the logical volumes. You can delete LUNs or CKD volumes, resize LUNs or volumes, and reuse the extents of those LUNs to create other LUNs or volumes, including ones of different sizes. One logical volume can be removed without affecting the other logical volumes that are defined on the same extent pool.

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

Two extent allocation methods (EAMs) are available for the DS8000: *Storage pool striping* (rotate extents) and *rotate volumes*.

Storage pool striping: Extent rotation

The storage allocation method (SAM) is chosen when a LUN or volume is created. The extents of a volume can be striped across several ranks. An extent pool with more than one rank is needed to use this SAM.

Note: Although the preferred SAM was storage pool striping, it is now a better choice to let Easy Tier manage the storage pool extents. This chapter describes rotate extents for the sake of completeness, but it is now mostly irrelevant.

The DS8900F maintains a sequence of ranks. The first rank in the list is randomly picked at each power-on of the storage system. The DS8900F tracks the rank in which the last allocation started. The allocation of the 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. Therefore, the system rotates the extents across the ranks, as shown in Figure 4-8.




Figure 4-8 Rotate extents

Rotate volumes allocation method

Extents can be allocated sequentially. In this case, all extents are taken from the same rank until enough extents are available for the requested volume size or the rank is full. In this 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 several volumes are provisioned (allocated), rotate through the ranks, as shown in Figure 4-9.

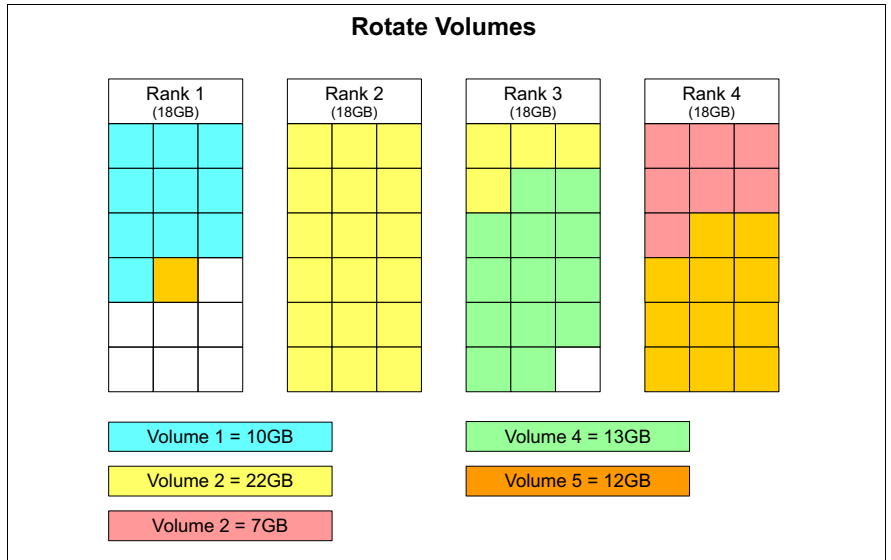


Figure 4-9 Rotate volumes

You might want to consider this allocation method if you prefer to manage performance manually. The workload of one volume is going to one rank. This configuration makes the identification of performance bottlenecks easier. However, by putting all the volumes' data onto one rank, you might introduce a bottleneck, depending on your actual workload.

Important: Rotate extents and rotate volume EAMs provide distribution of volumes over ranks. Rotate extents run this distribution at a granular (1 GiB or 16 MiB extent) level, and is a better method to minimize hot spots and improve overall performance.

However, as previously stated, Easy Tier is the preferred choice for managing the storage pool extents.

In a mixed-tier extent pool that contains different *tiers* of ranks, the storage pool striping EAM is used independently of the requested EAM, and the EAM is set to managed.

The Easy Tier default allocation order is High Performance. With the High Performance setting, the system populates drive classes in this order: Flash Tier 0, Flash Tier 1, and then Flash Tier 2.

There is a GUI and a CLI option for the whole Storage Facility to change the allocation preference. The two options are High Performance and High Utilization. With the High Utilization setting, the machine populates drive classes in this order: Flash Tier 1, Flash Tier 2, and then Flash Tier 0. The `chsi` command can be used to switch the `ETTierOrder` parameter between High Performance and High Utilization.

When you create striped volumes and non-striped volumes in an extent pool, a rank might be filled before the others. A full rank is skipped when you create striped volumes.

By using striped volumes, you distribute the I/O load of a LUN or CKD volume to more than one set of eight drives, which can enhance performance for a logical volume. In particular, OSs that do not include a volume manager with striping capability benefit most from this allocation method.

Small extents can increase the parallelism of sequential writes. Although the system stays within one rank until 1 GiB is written, with small extents it jumps to the next rank after 16 MiB. This configuration uses more disk drives when performing sequential writes.

Important: If you must add capacity to an extent pool because it is nearly full, it is better to add several ranks concurrently, not just one rank. This method allows new volumes to be striped across the newly added ranks.

With the Easy Tier manual mode facility, if the extent pool is a single-tier pool, the user can request an extent pool merge followed by a volume relocation with striping to run the same function. For a multitiered managed extent pool, extents are automatically relocated over time, according to performance needs. For more information, see *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667.

Rotate volume EAM: The rotate volume EAM is not allowed if one extent pool is composed of flash drives and configured for effective (virtual) capacity.

Dynamic Volume Expansion

The DS8900F supports expanding the size of a LUN or CKD volume without data loss by adding extents to the volume. The OS must support resizing.

A logical volume includes the attribute of being striped across the ranks or not. If the volume was created as striped across the ranks of the extent pool, the extents that are used to increase the size of the volume are striped. If a volume was created without striping, the system tries to allocate the additional extents within the same rank that the volume was created from originally.

Important: Before you can expand a volume, you must delete any CS relationship that involves that volume.

Because most OSs have no means of moving data from the end of the physical drive off to unused space at the beginning of the drive, and because of the risk of data corruption, IBM does not support shrinking a volume. The DS CLI and DS GUI interfaces cannot reduce the size of a volume.

Dynamic Volume Relocation

DVR is a capability that is provided as part of the Easy Tier manual mode facility.

DVR allows data that is stored on a logical volume to be migrated from its allocated storage to newly allocated storage while the logical volume remains accessible to attached hosts.

The user can request DVR by using the Migrate Volume function that is available through the DS GUI or DS CLI. DVR allows the user to specify a target extent pool and an EAM. The target extent pool can be a separate extent pool than the extent pool where the volume is. It can also be the same extent pool, but only if it is a single-tier pool. However, the target extent pool must be managed by the same DS8900F internal server.

Important: DVR in the same extent pool is not allowed in a managed pool. In managed extent pools, Easy Tier automatic mode automatically relocates extents within the ranks to allow performance rebalancing.

You can move volumes only among pools of the *same extent size*.

DVR provides the following capabilities:

- ▶ The ability to change the extent pool in which a logical volume is provisioned. This ability provides a mechanism to change the underlying storage characteristics of the logical volume to include the drive class (Tier 0, Tier 1, or Tier 2), and RAID array type. Volume migration can also be used to migrate a logical volume in to or out of an extent pool.
- ▶ The ability to specify the EAM for a volume migration that allows the EAM to be changed among the available EAMs anytime after volume creation. Volume migration that specifies the rotate extents EAM can also be used (in single-tier extent pools) to redistribute a logical volume's extent allocations across the existing ranks in the extent pool if more ranks are added to an extent pool.

Each logical volume has a configuration state. To begin a volume migration, the logical volume initially must be in the normal configuration state.

More functions are associated with volume migration that allow the user to pause, resume, or cancel a volume migration. Any or all logical volumes can be requested to be migrated at any time if available capacity is sufficient to support the reallocation of the migrating logical volumes in their specified target extent pool. For more information, see *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667.

4.4.4 Volume allocation and metadata

The DS8900F internal data layout is identical to the DS8880 internal data layout.

In earlier DS8000 systems, metadata allocation differed between fully provisioned or thin-provisioned (ESE) volumes:

- ▶ For fully provisioned volumes, metadata was contained in system-reserved areas of the physical arrays outside of the client data logical extent structure.

Figure 4-10 shows the extent layout in the DS8870 and earlier DS8000 models for standard, fully provisioned volumes. Note the size of the reserved area (not to scale), which includes volume metadata. The rest of the space is composed of 1 GB extents.

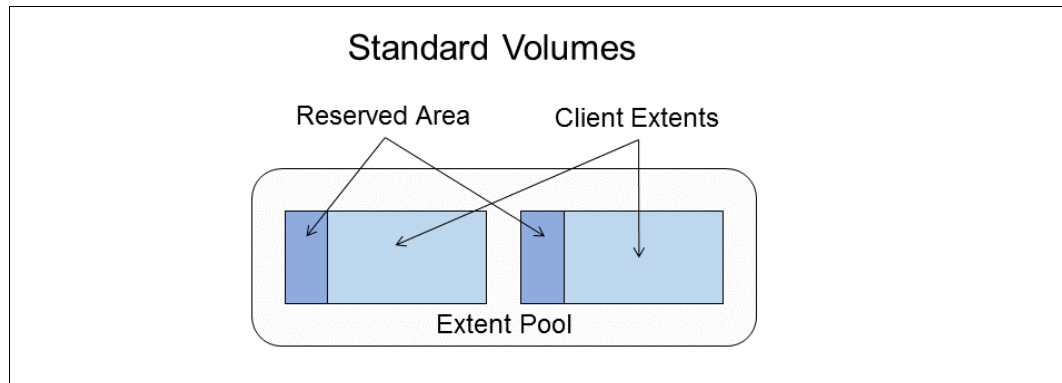


Figure 4-10 DS8000 standard volumes

- ▶ For ESE volumes, several logical extents, which are designated as *auxiliary rank extents*, are allocated to contain volume metadata.

Figure 4-11 shows the extent layout in the DS8870 and earlier for ESE volumes. In addition to the reserved area, auxiliary 1 GB extents are also allocated to store metadata for the ESE volumes.

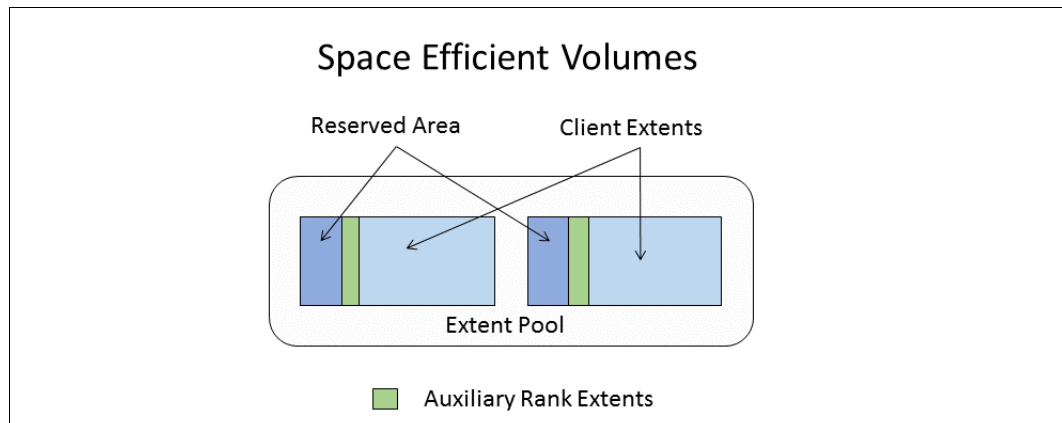


Figure 4-11 DS8870 extent space efficient volumes

In the DS8900F and DS8880, all volumes use an improved metadata extent structure, similar to what was previously used for ESE volumes. This unified extent structure greatly simplifies the internal management of logical volumes and their metadata. No area is fixed and reserved for volume metadata, and this capacity is added to the space that is available for use.

The metadata is allocated in the storage pool when volumes are created, and the space that is used by metadata is referred to as the *pool overhead*. Pool overhead means that the amount of space that can be provisioned (allocated) by volumes is variable and depends on both the number of volumes and the logical capacity of these volumes.

Tips: The following characteristics apply to the DS8900F:

- ▶ Metadata is variable in size. Metadata is used only as logical volumes are defined.
- ▶ The size of the metadata for a logical volume depends on the number and size of the logical volumes that are being defined.
- ▶ As volumes are defined, the metadata is immediately allocated and reduces the size of the available user extents in the storage pool.
- ▶ As you define new logical volumes, you must allow for the fact that this metadata space continues to use user extent space.

For storage pools with large extents, metadata is also allocated as large extents (1 GiB for FB pools or 1113 cylinders for CKD pools). Large extents that are allocated for metadata are subdivided into 16 MiB subextents, which are also referred to as *metadata extents*, for FB volumes, or 21 cylinders for CKD. For extent pools with small extents, metadata extents are also small extents. Sixty-four metadata subextents are in each large metadata extent for FB, and 53 metadata subextents are in each large metadata extent for CKD.

For each FB volume that is provisioned (allocated), an initial 16 MiB metadata subextent or metadata small extent is allocated, and an extra 16 MiB metadata subextent or metadata small extent is allocated for every 10 GiB of provisioned (allocated) capacity or portion of provisioned capacity.

For each CKD volume that is provisioned (allocated), an initial 21 cylinders metadata subextent or metadata small extent is allocated, and an extra 21 cylinders metadata subextent or metadata small extent is allocated for every 11130 cylinders (or ten 3390 Model 1) of allocated capacity or portion of allocated capacity.

For example, a 3390-3 (that is, 3339 cylinders or about 3 GB) or 3390-9 (that is, 10,017 cylinders or 10 GB) volume takes two metadata extents (one metadata extent for the volume and another metadata extent for any portion of the first 10 GB). A 128 GB FB volume takes 14 metadata extents (one metadata extent for the volume and another 13 metadata extents to account for the 128 GB).

Figure 4-12 on page 127 shows an illustration of 3 GB and 12 GB FB volumes for a storage pool with large extents. In an extent pool with small extents, there is no concept of subextents. You have user extents, unused extents, and metadata extents.

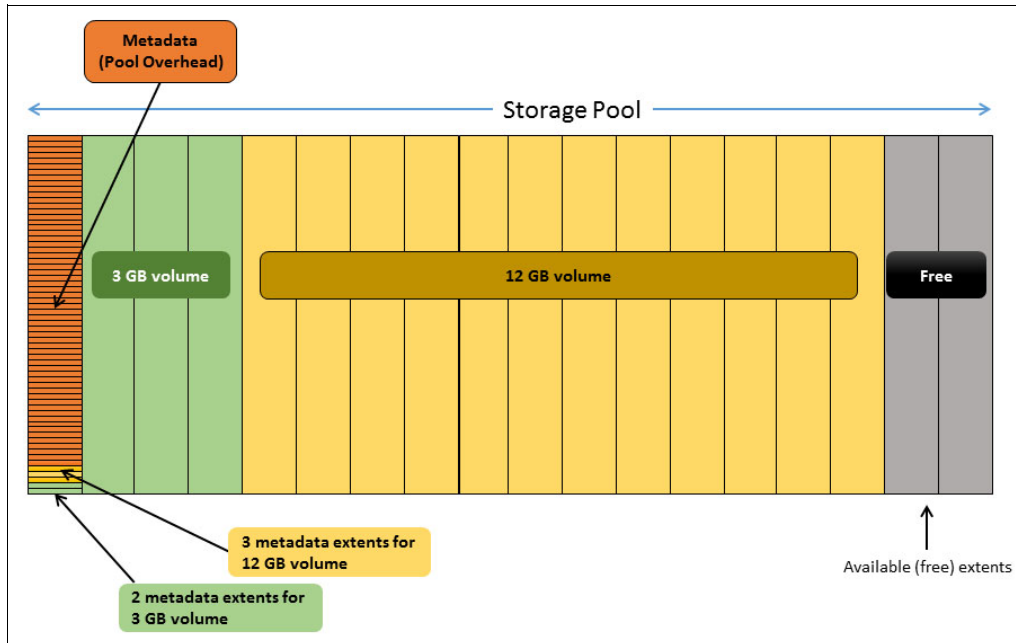


Figure 4-12 Metadata allocation

Metadata extents with free space can be used for metadata by any volume in the extent pool.

When you use metadata extents and user extents within an extent pool, some planning and calculations are required, especially in a mainframe environment where thousands of volumes are defined and the whole capacity is provisioned (allocated) during the initial configuration. You must calculate the capacity that is used up by the metadata to get the capacity that can be used for user data. This calculation is important only when fully provisioned volumes are used. Thin-provisioned volumes use no space when created; only metadata and space are used when data is written.

For extent pools with small extents, the number of available user data extents can be estimated as follows:

$$(user\ extents) = (pool\ extents) - (number\ of\ volumes) - (total\ provisioned\ capacity)/10$$

For extent pools with regular 1 GiB extents where the details of the volume configuration are not known, you can estimate the number of metadata extents based on many volumes only. The calculations are performed as shown:

- ▶ FB pool overhead = $(number\ of\ volumes \times 2 + total\ volume\ extents/10)/64$ and rounded up to the nearest integer
- ▶ CKD pool overhead = $(number\ of\ volumes \times 2 + total\ volume\ extents/10)/53$ and rounded up to the nearest integer

The formulas overestimate the space that is used by the metadata by a small amount because it assumes wasted space on every volume. However, the precise size of each volume does not need to be known.

Here are two examples:

- ▶ A CKD storage pool has 6,190 extents in which you expect to allocate all of the space on 700 volumes. Your pool overhead is 39 extents by using this calculation:

$$(700 \times 2 + 6190/10)/53 = 38.09$$

- ▶ An FB storage pool has 6,190 extents in which you expect to use thin provisioning and allocate up to 12,380 extents (2:1 overprovisioning) on 100 volumes. Your pool overhead is 23 extents by using this calculation:

$$(100 \times 2 + 12380/10)/64 = 22.46$$

Total volume effective (virtual) capacity

IBM is standardizing the capacity names. A list of standardized names is listed in 4.5, “Terminology for IBM Storage products” on page 137.

Space for Easy Tier

The Easy Tier function of DS8000 monitors access to extents to host applications. Easy Tier moves extents to higher or lower tiers depending on the access frequency or *heat* of an extent. To move around extents, some free space or free extents must be available in an extent pool. You can manually watch the free space or extents, but a better option is to let the system reserve some space for Easy Tier extent movements. Run the `chsi` command to change the `ETSMode` parameter to `Enabled`. This setting causes the storage system to reserve some capacity for Easy Tier. Obviously, when this mode is enabled, capacity for new volume allocations is reduced.

Extent space efficient volumes

Volumes can be created as thin-provisioned or fully provisioned.

Space for a thin-provisioned volume is allocated when a write occurs. More precisely, it is allocated when a destage from the cache occurs and insufficient free space is left on the currently allocated extent.

Therefore, thin provisioning allows a volume to exist that is larger than the usable (physical) capacity in the extent pool to which it belongs. This approach allows the “host” to work with the volume at its defined capacity, even though insufficient usable (physical) space might exist to fill the volume with data.

The assumption is that either the volume is never filled, or as the DS8900F runs low on raw capacity, more is added. This approach also assumes that the DS8900F is not at its maximum raw capacity.

Thin-provisioned volume support is contained in the BF license group. This provisioning is an attribute that you specify when creating a volume.

Note: If thin provisioning is used, the metadata is allocated for the entire volume (effective provisioned capacity) when the volume is created, not when extents are used.

CKD thin provisioning

Thin provisioning is also possible with CKD volumes. See *IBM DS8880 Thin Provisioning (Updated for Release 8.5)*, REDP-5343, and check for the APARs that are needed for certain z/OS versions.

Extent space efficient capacity controls for thin provisioning

Using thin provisioning can affect the amount of storage capacity that you choose to order. Use ESE capacity controls to allocate storage correctly.

With the mixture of thin-provisioned (ESE) and fully provisioned (non-ESE) volumes in an extent pool, a method is needed to dedicate part of the extent-pool storage capacity for ESE user data usage, and to limit the ESE user data usage within the extent pool.

Also, you must be able 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 that is available for ESE user data usage. These controls also ensure that a proportion of the extent pool storage is available for ESE user data usage.

The following controls are available to limit the usage of extents in an extent pool:

- ▶ Reserve capacity in an extent pool by enabling the extent pool limit function by running the `chextpool -extentlimit enable -limit extent_limit_percentage pool_ID` command.
- ▶ You can reserve space for the sole use of ESE volumes by creating a repository by running the `mksestg -repcap capacity pool_id` command.

Figure 4-13 shows the areas in an extent pool.




Figure 4-13 Spaces in an extent pool

Capacity controls exist for an extent pool and also for a repository, if it is defined. There are system-defined warning thresholds at 15% and 0% free capacity left, and you can set your own user-defined threshold for the whole extent pool or for the ESE repository. Thresholds for an *extent pool* are set by running the DS CLI `chextpool -threshold` (or `mkextpool`) command. Thresholds for a *repository* are set by running the `chsestg -repcapthreshold` (or `mksestg`) command.

The `threshold exceeded` status refers to the user-defined threshold.

The extent pool or repository status shows one of the three following values:

- ▶ 0: The percent of available capacity is greater than the extent pool or repository threshold.
- ▶ 1: The percent of available capacity is greater than zero but less than or equal to the extent pool or repository threshold.
- ▶ 10: The percent of available capacity is zero.

A Simple Network Management Protocol (SNMP) trap is associated with the extent pool / repository capacity controls, and it notifies you when the extent usage in the pool exceeds a user-defined threshold and when the extent pool is out of extents for user data.

When the size of the extent pool remains fixed or when it increases, the available usable (physical) capacity remains greater than or equal to the *provisioned (allocated)* capacity. However, a reduction in the size of the extent pool can cause the *available* usable (physical) capacity to become less than the *provisioned (allocated)* capacity in certain cases.

For example, if the user requests that one of the ranks in an extent pool is depopulated, the data on that rank is moved to the remaining ranks in the pool. This process causes the rank to become not provisioned (allocated) and removed from the pool. The user is advised to inspect the limits and threshold on the extent pool after any changes to the size of the extent pool to ensure that the specified values are still consistent with the user's intentions.

Overprovisioning control

It is possible to set the maximum allowed overprovisioning ratios for an extent pool.

A new attribute (**-opratiolimit**) is available when creating or modifying extent pools to add operational limits. Example 4-3 provides an example of creating and modifying an extent pool with a defined operational limit that cannot be exceeded.

Example 4-3 The opratiolimit parameter

```
dsscli> mkextpool -rankgrp 0 -stgtype fb -opratiolimit 3.5 -encryptgrp 1 my_pool
CMUC00000I mkextpool: Extent pool P3 successfully created.
dsscli> chextpool -opratiolimit 3.125 p3
CMUC00001I chextpool: Extent pool P3 successfully modified.
```

Setting an overprovisioning ratio limit results in the following changes to system behavior to prevent an extent pool from exceeding the overprovisioning ratio:

- ▶ Prevent volume creation, expansion, or migration.
- ▶ Prevent rank depopulation.
- ▶ Prevent pool merge.
- ▶ Prevent turning on Easy Tier space reservation.

For more information, see *IBM DS8880 Thin Provisioning (Updated for Release 8.5)*, REDP-5343.

4.4.5 Logical subsystems

An LSS is another logical construct. It can also be referred to as an LCU. The term *LSS* is mostly used in association with FB volumes. The term *LCU* is used in association with CKD volumes. A maximum of 255 LSSs can exist in the DS8000. Each of them have an identifier of 00 - FE. An individual LSS must manage either FB or CKD volumes.

All even-numbered LSSs (X'00', X'02', X'04', up to X'FE') are handled by internal server 0, and all odd-numbered LSSs (X'01', X'03', X'05', up to X'FD') are handled by internal server 1. LSS X'FF' is reserved. This configuration allows both servers to handle host commands to the volumes in the DS8000 if the configuration takes advantage of this capability. If either server is not available, the remaining operational server handles all LSSs. LSSs are also placed in address groups of 16 LSSs, except for the last group that has 15 LSSs. The first address group is 00 - 0F, and so on, until the last group, which is F0 - FE.

Because the LSSs manage volumes, an individual LSS must manage the same type of volumes. An address group must also manage the same type of volumes. The first volume (either FB or CKD) that is assigned to an LSS in any address group sets that group to manage those types of volumes. For more information, see "Address groups" on page 132.

Volumes are created in extent pools that are associated with either internal server 0 or 1. Extent pools are also formatted to support either FB or CKD volumes. Therefore, volumes in any internal server 0 extent pools can be managed by any even-numbered LSS if the LSS and extent pool match the volume type. Volumes in any internal server 1 extent pools can be managed by any odd-numbered LSS if the LSS and extent pool match the volume type.

Volumes also have an identifier 00 - FF. The first volume that is assigned to an LSS has an identifier of 00. The second volume is 01, and so on, up to FF if 256 volumes are assigned to the LSS.

For FB volumes, the LSSs that are used to manage them are not significant if you spread the volumes between odd and even LSSs. When the volume is assigned to a host (in the DS8900F configuration), a LUN is assigned to it that includes the LSS and Volume ID. This LUN is sent to the host when it first communicates with the DS8900, so it can include the LUN in the "frame" that is sent to the DS8900F when it wants to run an I/O operation on the volume. This method is how the DS8900F knows which volume on which to run the operation.

Conversely, for CKD volumes, the LCU *is* significant. The LCU must be defined in a configuration that is called the input/output configuration data set (IOCDS) on the host. The LCU definition includes a control unit address (CUADD). This CUADD must match the LCU ID in the DS8900F. A device definition for each volume, which has a unit address (UA) that is included, is also included in the IOCDS. This UA must match the volume ID of the device. The host must include the CUADD and UA in the "frame" that is sent to the DS8900F when it wants to run an I/O operation on the volume. This method is how the DS8900F knows which volume on which to run the operation.

For both FB and CKD volumes, when the "frame" that is sent from the host arrives at a host adapter port in the DS8900F, the adapter checks the LSS or LCU identifier to know which internal server to pass the request to inside the DS8900. For more information about host access to volumes, see 4.4.6, "Volume access" on page 133.

FB LSSs are created automatically when the first FB logical volume on the LSS is created. FB LSSs are deleted automatically when the last FB logical volume on the LSS is deleted. CKD LCUs require user parameters to be specified and must be created before the first CKD logical volume can be created on the LCU. They must be deleted manually after the last CKD logical volume on the LCU is deleted.

Certain management actions in Metro Mirror (MM), Global Mirror (GM), or Global Copy (GC) operate at the LSS level. For example, the freezing of pairs to preserve data consistency across all pairs in case a problem occurs with one of the pairs is performed at the LSS level.

The option to put all or most of the volumes of a certain application in one LSS makes the management of remote copy operations easier, as shown in Figure 4-14.

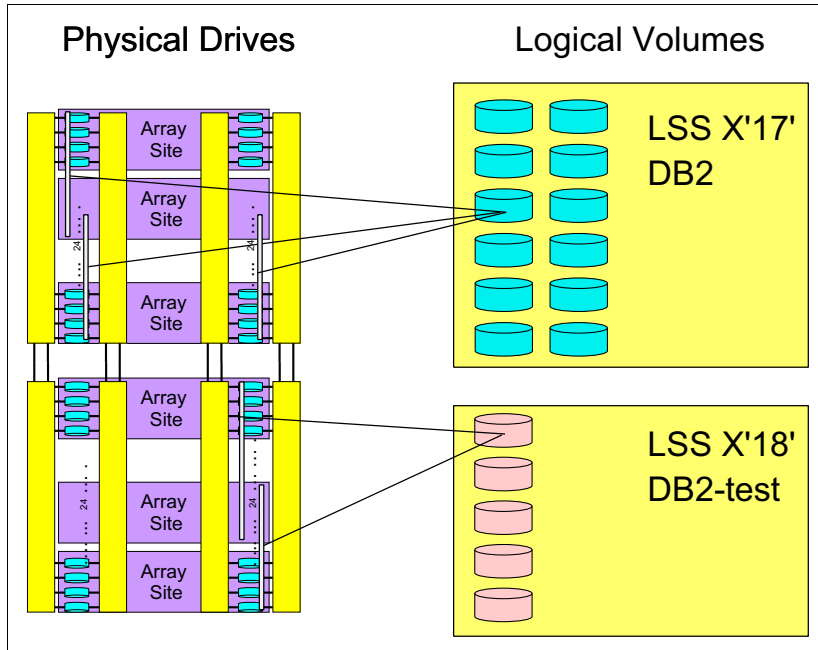


Figure 4-14 Grouping of volumes in LSSs

Address groups

Address groups are created automatically when the first LSS that is associated with the address group is created. The groups are deleted automatically when the last LSS in the address group is deleted.

All devices in an LSS must be CKD or FB. This restriction goes even further. LSSs are grouped into address groups of 16 LSSs. LSSs are numbered X'ab', where *a* is the address group and *b* denotes an LSS within the address group. For example, X'10' - X'1F' are LSSs in address group 1.

All LSSs within one address group must be of the same type (CKD or FB). The first LSS that is defined in an address group sets the type of that address group.

Figure 4-15 shows the concept of LSSs and address groups.

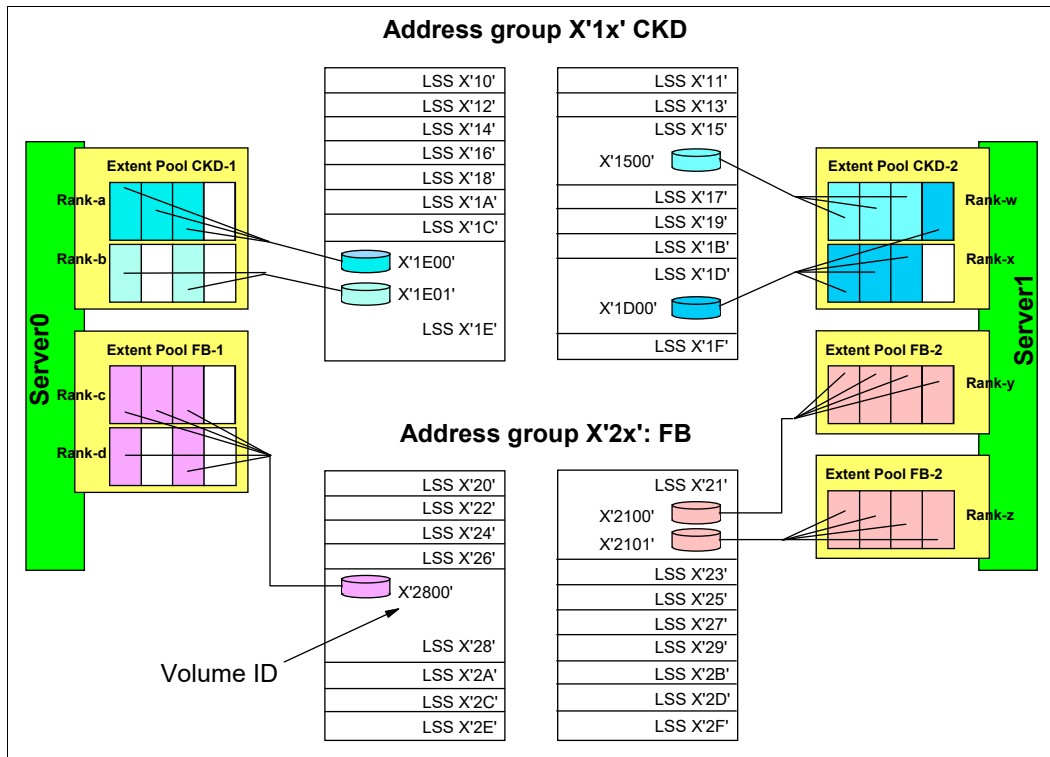


Figure 4-15 Logical storage subsystems

The LUN identification $X'gabb'$ is composed of the address group $X'g'$, the LSS number within the address group $X'a'$, and the ID of the LUN within the LSS $X'bb'$. For example, FB LUN $X'2101'$ denotes the second ($X'01'$) LUN in LSS $X'21'$ of address group 2.

An extent pool can have volumes that are managed by multiple address groups. The example in Figure 4-15 shows only one address group that is used with each extent pool.

4.4.6 Volume access

The DS8900F provides mechanisms to control host access to LUNs. In most cases, a host system features two or more HBAs and the system needs access to a group of LUNs. For easy management of system access to logical volumes, the DS8900F introduces the concept of host attachments and volume groups.

Host attachment

HBAs are identified to the DS8900F in a host attachment construct that specifies the worldwide port names (WWPNs) of a host's HBAs. A set of host ports can be associated through a port group attribute that allows a set of HBAs to be managed collectively. This port group is referred to as a *host attachment* within the configuration.

Each host attachment can be associated with a volume group to define the LUNs that host is allowed to access. Multiple host attachments can share the volume group. The host attachment can also specify a port mask that controls the DS8900F I/O ports that the host HBA is allowed to log in to. Whichever ports the HBA logs in to, it sees the same volume group that is defined on the host attachment that is associated with this HBA.

The maximum number of host attachments on a DS8900F is 8,192. This host definition is required only for open systems hosts. Any IBM Z server can access any CKD volume in a DS8000 if its IOCDs is correct.

Volume group

A *volume group* is a named construct that defines a set of logical volumes. A volume group is required only for FB volumes. When a volume group is used with CKD hosts, a default volume group contains all CKD volumes. Any CKD host that logs in to a Fibre Channel connection (IBM FICON) I/O port has access to the volumes in this volume group. CKD logical volumes are automatically added to this volume group when they are created and are automatically removed from this volume group when they are deleted.

When a host attachment object is used with open systems hosts, a host attachment object that identifies the HBA is linked to a specific volume group. You must define the volume group by indicating the FB volumes that are to be placed in the volume group. Logical volumes can be added to or removed from any volume group dynamically.

Important: Volume group management is available only with the DS CLI. In the DS GUI, users define hosts and assign volumes to hosts. A volume group is defined in the background. No volume group object can be defined in the DS GUI.

Two types of volume groups are used with open systems hosts. The type determines how the logical volume number is converted to a host addressable LUN_ID in the Fibre Channel (FC) SCSI interface. A *SCSI map volume group* type is used with FC SCSI host types that poll for LUNs by walking the address range on the SCSI interface. This type of volume group can map any FB logical volume numbers to 256 LUN IDs that have zeros in the last 6 bytes and the first 2 bytes in X'0000' - X'00FF'.

A *SCSI mask volume group* type is used with FC SCSI host types that use the **Report LUNs** command to determine the LUN IDs that are accessible. This type of volume group can allow any FB logical volume numbers to be accessed by the host where the mask is a bitmap that specifies the LUNs that are accessible. For this volume group type, the logical volume number X'abcd' is mapped to LUN_ID X'40ab40cd00000000'. The volume group type also controls whether 512-byte block LUNs or 520-byte block LUNs can be configured in the volume group.

When a host attachment is associated with a volume group, the host attachment contains attributes that define the logical block size and the Address Discovery Method (LUN Polling or Report LUNs) that is used by the host HBA. These attributes must be consistent with the volume group type of the volume group that is assigned to the host attachment. This consistency ensures that HBAs that share a volume group have a consistent interpretation of the volume group definition and have access to a consistent set of logical volume types.

The DS Storage Manager GUI typically sets these values for the HBA based on your specification of a host type. You must consider what volume group type to create when a volume group is set up for a particular HBA.

FB logical volumes can be defined in one or more volume groups. This definition allows a LUN to be shared by host HBAs that are configured to separate volume groups. An FB logical volume is automatically removed from all volume groups when it is deleted.

The DS8900F supports a maximum of 8,320 volume groups.

Figure 4-16 shows the relationships between host attachments and volume groups. Host AIXprod1 has two HBAs, which are grouped in one host attachment, and both HBAs are granted access to volume group DB2-1. Most of the volumes in volume group DB2-1 are also in volume group DB2-2, which is accessed by the system AIXprod2.

However, one volume in each group is not shared in Figure 4-16. The system in the lower-left part of the figure has four HBAs, and they are divided into two distinct host attachments. One HBA can access volumes that are shared with AIXprod1 and AIXprod2. The other HBAs have access to a volume group that is called docs.

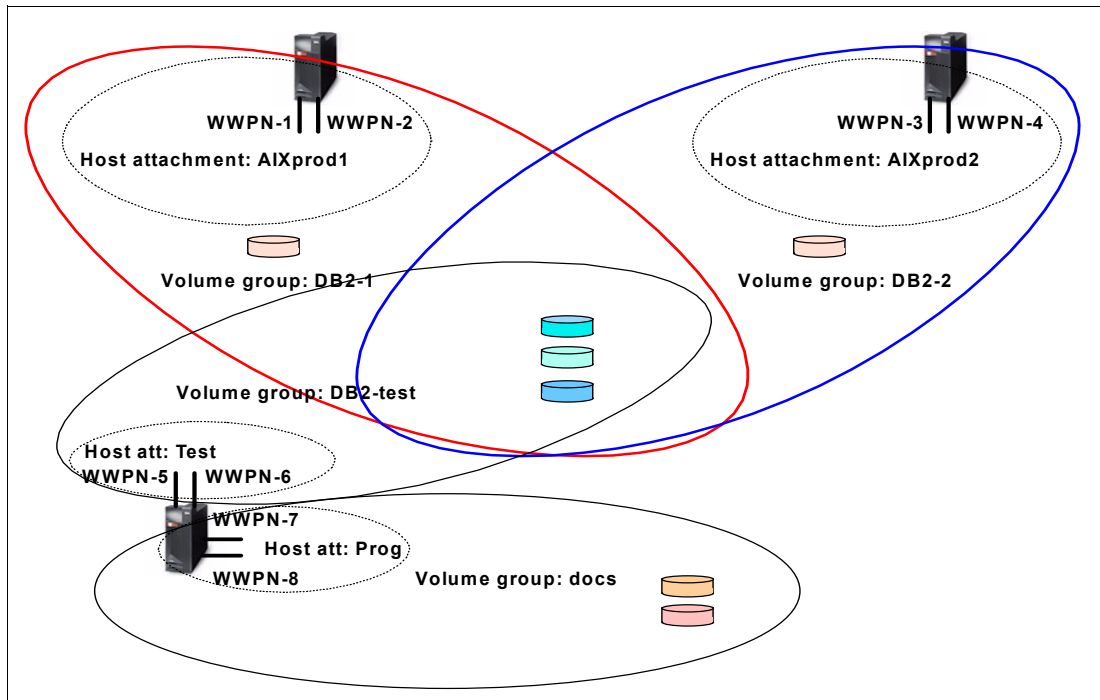


Figure 4-16 Host attachments and volume groups

When working with Open Systems clusters for defining volumes, use the Create Cluster function of the DS Storage Manager GUI to easily define volumes. In general, the GUI hides the complexity of certain DS8000 internal definition levels like volume groups, array sites, and ranks. It helps save time by directly processing these definitions internally in the background without presenting them to the administrator.

4.4.7 Virtualization hierarchy summary

As you view the virtualization hierarchy (as shown in Figure 4-17 on page 136), start with many drives that are grouped in array sites. The array sites are created automatically when the drives are installed.

Complete the following steps as a user:

1. Associate an array site with a RAID array, and allocate spare drives as required.
2. Associate the array with a rank with small or large extents that are formatted for FB or CKD data.

3. Add the extents from the selected ranks to an extent pool. The combined extents from the ranks in the extent pool are used for subsequent allocation for one or more logical volumes. Within the extent pool, you can reserve space for ESE volumes. ESE volumes require effective (virtual) capacity to be available in the extent pool.
4. Create logical volumes within the extent pools (by default, striping the volumes), and assign them a logical volume number that determines the LSS that they are associated with and the internal server that manages them. The LUNs are assigned to one or more volume groups.
5. Configure the host HBAs into a host attachment that is associated with a volume group.

Figure 4-17 shows the virtualization hierarchy.

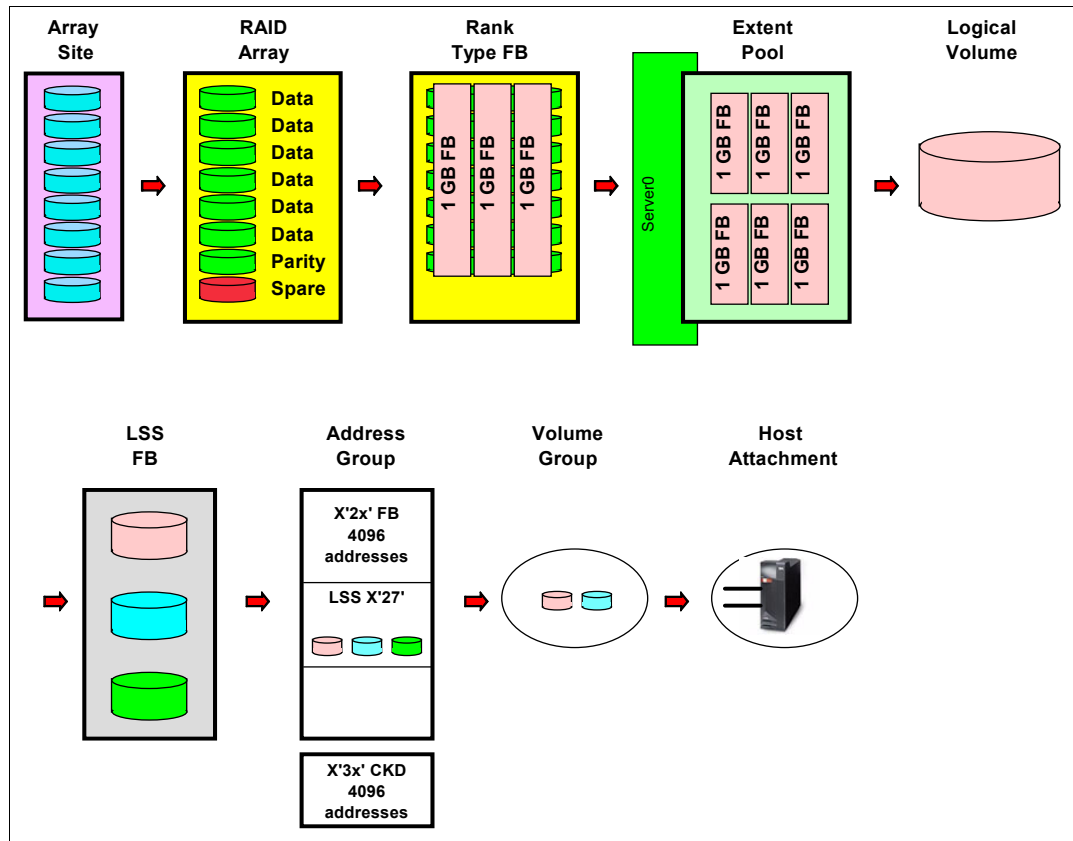


Figure 4-17 Virtualization hierarchy

This virtualization concept provides a high degree of flexibility. Logical volumes can be dynamically created, deleted, and resized. They can also be grouped logically to simplify storage management. Large LUNs and CKD volumes reduce the total number of volumes, which contributes to the reduction of management effort.

Tip: The DS GUI helps save administration steps by handling some of these virtualization levels in the background and automatically processing them for the administrator.

4.5 Terminology for IBM Storage products

This section lists the capacity terms for all IBM Storage products. These terms replace all terms that were previously used to describe capacity in IBM Storage products.

▶ Raw capacity

The reported capacity of the drives in the system before formatting or RAID is applied.

▶ Usable capacity

The amount of capacity that is available for storing data on a system, pool, or array after formatting and RAID techniques are applied.

▶ Used (usable) capacity

The amount of usable capacity that is taken up by data in a system, pool, or array after data reduction techniques are applied.

▶ Available (usable) capacity

The amount of usable capacity that is not yet used in a system, pool, or array.

▶ Provisioned capacity

The amount of provisioned capacity that can be created in a system or pool without running out of usable capacity for the current data-reduction savings that you want achieved. This capacity equals the usable capacity that is divided by the data reduction savings percentage. In some storage systems, restrictions in the system determine the maximum provisioned capacity that is allowed in a pool or system. In those cases, the provisioned capacity cannot exceed this limit.

The current implementation of the DS8900F does not feature any data reduction techniques.

Provisioned capacity was previously called *virtual capacity* in this chapter.

▶ Written capacity

The amount of usable capacity that would have been used to store written data in a pool or system if data reduction was not applied.

▶ Overhead capacity

The amount of usable capacity that is occupied by metadata in a pool or system and other data that is used for system operation.

▶ Thin-provisioning savings

The total amount of usable capacity that is saved in a pool, system, or volume by using usable capacity when needed as a result of write operations. The capacity that is saved is the difference between the provisioned capacity minus the written capacity.

▶ Overprovisioning

The result of creating more provisioned capacity in a storage system or pool than there is usable capacity. This result occurs when thin provisioning or compression ensures that the used capacity of the provisioned volumes is less than their provisioned capacity.

▶ Overprovisioned ratio

The ratio of provisioned capacity to usable capacity in a pool or system.



Part 2

Planning and installation

This part describes the installation planning process for the IBM DS8900.

This part contains the following chapters:

- ▶ Chapter 5, “IBM DS8900F physical planning and installation” on page 141
- ▶ Chapter 6, “IBM DS8900F Management Console planning and setup” on page 167
- ▶ Chapter 7, “IBM DS8900F features and licensed functions” on page 199



IBM DS8900F physical planning and installation

This chapter describes the various steps that are involved in the planning and installation of the IBM DS8900F. It includes a reference listing of the information that is required for the setup and where to find detailed technical reference material.

This chapter covers the following topics:

- ▶ Considerations before the installation: Planning for growth
- ▶ Planning for the physical installation
- ▶ Network connectivity planning
- ▶ Remote Mirror and Remote Copy connectivity
- ▶ Disk capacity considerations

For more information about the configuration and installation process, see the *IBM DS8900F Introduction and Planning Guide*, SC27-9560.

5.1 Considerations before the installation: Planning for growth

Start by developing and following a project plan to address the necessary topics for a successful implementation. Consider the following items for your installation plan checklist:

- ▶ Plan for growth to minimize disruption to operations.

Important: The IBM DS8980F and DS8950F systems support an expansion frame that can be installed adjacent or 20 meters away from the base frame. (Feature Code 1341 is needed.)

- ▶ Consider location suitability, floor loading, access constraints, elevators, and doorways.
- ▶ Analyze power requirements, such as redundancy and using an uninterruptible power supply (UPS).
- ▶ Examine environmental requirements, such as adequate cooling capacity.
- ▶ Full Disk Encryption (FDE) drives are a standard feature for the DS8900F. If encryption activation is required, consider the location and connection needs for the external key servers, such as IBM Security Key Lifecycle Manager or Gemalto SafeNet KeySecure servers.
- ▶ Consider the integration of Lightweight Directory Access Protocol (LDAP) to allow a single user ID and password management. LDAP can be configured from the Storage Management GUI, as described in 6.5.2, “Remote authentication” on page 188.
- ▶ Call Home through a Secure Sockets Layer (SSL) installation to provide a continued secure connection to the IBM Support center.
- ▶ Consider connecting to IBM Storage Insights that can help you predict and prevent storage problems before they impact your business.
- ▶ Plan for logical configuration, Copy Services (CS), and staff education. For more information, see Chapter 8, “Configuration flow” on page 225.

5.1.1 Client responsibilities for the installation

The DS8900F is specified as an IBM or IBM Business Partner installation and setup system. However, the following activities are several required planning and installation activities for which the client is responsible at a high level:

- ▶ Physical configuration planning
Your Storage Marketing Specialist can help you plan and select the DS8900F model physical configuration and features.
- ▶ Installation planning
- ▶ Integration of LDAP
IBM can help in planning and implementation at the client’s request.
- ▶ Installation of Assist On-site (AOS)
IBM can help plan and implement at the client’s request.
- ▶ Integration of IBM Spectrum Control and Simple Network Management Protocol (SNMP) into the client environment for monitoring of performance and configuration
IBM can provide services to set up and integrate these components.

- ▶ Configuration and integration of external key servers and IBM DS8000 Encryption for enhanced data security

Supported key servers for data at rest and Transparent Cloud Tiering (TCT) Encryption include IBM Security Key Lifecycle Manager, Gemalto Safenet KeySecure, and Thales Vormetric DSM. IBM Security Guardium Key Lifecycle Manager is the only supported key server for encryption of data in flight (IBM Fibre Channel Endpoint Security).

IBM provides services to set up and integrate IBM Security Guardium Key Lifecycle Manager components.

Alternatively, clients can install the Gemalto SafeNet key servers or Thales Vormetric DSM. For IBM Fibre Channel Endpoint Security, IBM Security Guardium Key Lifecycle Manager with Key Management Interoperability Protocol (KMIP) in Multi-Master mode is required.

- ▶ Logical configuration planning and application

Logical configuration refers to the creation of redundant array of independent disks (RAID) arrays and pools, and to the assignment of the configured capacity to servers. Application of the initial logical configuration and all subsequent modifications to the logical configuration also are client responsibilities. The logical configuration can be created, applied, and modified by using the DS GUI, DS Command-line Interface (DS CLI), or DS Open application programming interface (DS Open API).

IBM Services® also can apply or modify your logical configuration, which is a fee-based service.

5.1.2 Participants

A project manager must coordinate the many tasks that are necessary for a successful installation. Installation requires close cooperation with the user community, IT support staff, and technical resources that are responsible for floor space, power, and cooling.

A storage administrator must also coordinate requirements from the user applications and systems to build a storage plan for the installation. This plan is needed to configure the storage after the initial hardware installation is complete.

The following people must be briefed and engaged in the planning process for the physical installation:

- ▶ Systems and storage administrators
- ▶ Installation planning engineer
- ▶ Building engineer for floor loading, air conditioning, and electrical considerations
- ▶ Security engineers for AOS, LDAP, key servers, and encryption
- ▶ Administrator and operator for monitoring and handling considerations
- ▶ IBM Systems Service Representative (IBM SSR) or IBM Business Partner

5.1.3 Required information

A validation list to help the installation process must include the following items:

- ▶ Drawings that detail the DS8000 placement as specified and agreed upon with a building engineer, which ensures that the weight is within limits for the route to the final installation position.
- ▶ Approval to use elevators if the DS8900F weight and size are acceptable.
- ▶ Connectivity information, servers, storage area network (SAN), and mandatory local area network (LAN) connections.

- ▶ Agreement on the security structure of the installed DS8000 with all security engineers.
- ▶ Agreement on the detailed storage plan. Ensure that the configuration specialist has all of the information to configure all of the storage and set up the environment, as required.
- ▶ Activation codes for Base Functions (BFs), which are mandatory, and any optional feature activation codes.

5.2 Planning for the physical installation

This section describes the physical installation planning process and provides important tips and considerations.

5.2.1 Delivery and staging area

The shipping carrier is responsible for delivering and unloading the DS8900F as close to its final destination as possible. Inform the carrier of the weight and size of the packages to deliver. Also, inspect the site and the areas through which the packages will be moved (for example, hallways, floor protection, elevator size, and loading).

Table 5-1 lists the final packaged dimensions and maximum packaged weight of the DS8900F storage unit ship group. The maximum packaged weight is the maximum weight of the frame plus the packaging weight.

Table 5-1 Packaged dimensions and weight for DS8900F models

Shipping container	Packaged dimensions (in centimeters and inches)	Maximum packaged weight (in kilograms and pounds)
IBM DS8910F model 993	Height 1.49 m (58.7 in.) Width 1.05 m (41.3 in.) Depth 1.30 m (51.2 in.)	295 kg (650 lb)
DS8910F model 994	Height 2.22 m (87.7 in.) Width 0.95 m (37.4 in.) Depth 1.50 m (59.1 in.)	762 kg (1680 lb)
DS8950F model 996 and DS8980F model 998	Height 2.22 m (87.7 in.) Width 1.0 m (39.4 in.) Depth 1.50 m (59.1 in.)	793 kg (1748 lb)
Expansion Frame model E96	Height 2.22 m (87.7 in.) Width 1.0 m (39.4 in.) Depth 1.50 m (59.1 in.)	603 kg (1330 lb)

By using the shipping weight reduction option, you can receive delivery of a DS8900F model in multiple shipments that do not exceed 909 kg (2,000 lb) each.

The DS8910F model 993 can be integrated into an existing IBM z15 model T02, IBM LinuxONE Rockhopper III Model LT2, IBM z14 Model ZR1, IBM LinuxONE Rockhopper II Model LR1, or other standard 19-inch wide frame with 16U contiguous space. For more information, see *IBM DS8910F Model 993 Rack-Mounted Storage System Release 9.1*, REDP-5566.

For more information about the Shipping Weight Reduction option, see Chapter 7, “IBM DS8900F features and licensed functions” on page 199.

5.2.2 Floor type and loading

The DS8900F can be installed on a raised or nonraised floor. The total weight and space requirements of the storage unit depend on the configuration features that you ordered. You might consider calculating the weight of the unit and the expansion frame (if ordered) in their maximum capacity to allow for the addition of new features.

For the maximum weight of the various DS8900F models, see Table 5-1 on page 144.

Important: You must check with the building engineer or other appropriate personnel to ensure that the floor loading is correctly considered.

Figure 5-1 shows the location of the cable cutouts for DS8900F. You can use the following measurements when you cut the floor tile:

- ▶ Width: 41.91 cm (16.5 in.)
- ▶ Depth: 8.89 cm (3.5 in.)
- ▶ End of frame to edge of cable cutout: 10.0 cm (3.9 in.)




Figure 5-1 Floor tile cable cutout for DS8900F

For more information about floor loading and weight distribution, see *IBM DS8900F Introduction and Planning Guide*, SC27-9560.

5.2.3 Overhead cabling features

The overhead cabling (top exit) feature, as shown in Figure 5-2, is available for DS8900F as an alternative to the standard rear cable exit. Verify whether you ordered the top exit feature before the tiles for a raised floor are cut.

This feature requires the following items:

- ▶ Feature Code 1401 Top exit bracket for overhead cabling
- ▶ Feature Code 1101 Overhead ladder

For more information, see *IBM DS8900F Introduction and Planning Guide*, SC27-9560.

Figure 5-2 shows the overhead cabling (top exit) feature.

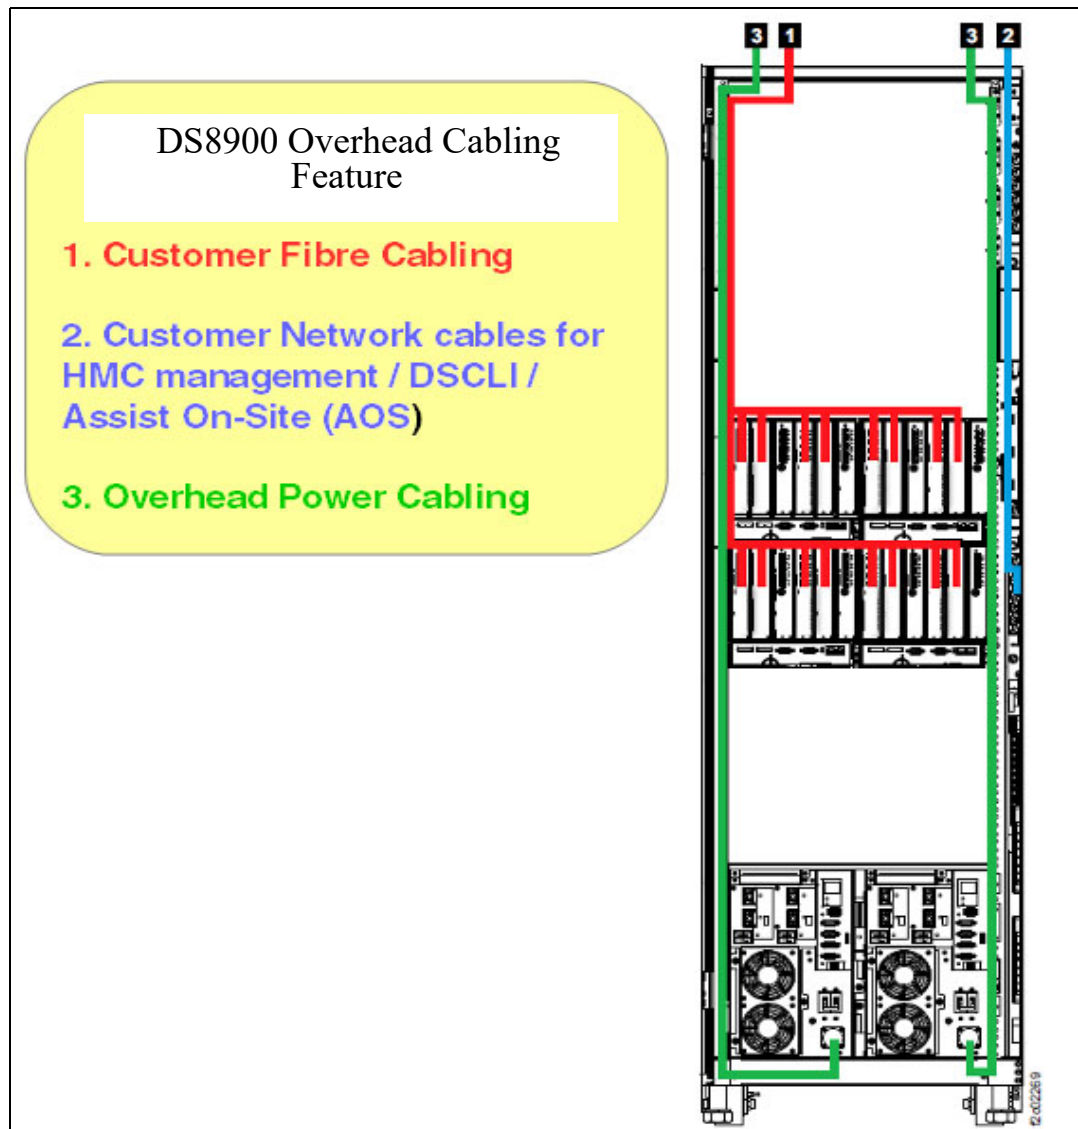


Figure 5-2 Overhead cabling for DS8900F

5.2.4 Room space and service clearance

The total amount of space that is needed by the storage units can be calculated by using the dimensions that are shown in Table 5-2.

Table 5-2 DS8900F dimensions

Dimensions with casters and covers	DS8900F all models (racked)
Height	193 cm (76 in.)
Width	64 cm (25 in.)
Depth	144 cm (56.5 in.)

The storage unit location area also covers the service clearance that is needed by the IBM SSR when the front and rear of the storage unit are accessed. You can use the following minimum service clearances. Verify your configuration and the maximum configuration for your needs, keeping in mind that the DS8900F has a maximum of one expansion frame (for a total of two frames).

Figure 5-3 shows the overall physical footprint of a storage system.




Figure 5-3 DS8900F service clearance requirements

The following clearances are needed:

- ▶ For the front of the unit, allow a minimum of 121.92 cm (48 in.).
- ▶ For the rear of the unit, allow a minimum of 76.2 cm (30 in.).

5.2.5 Power requirements and operating environment

Consider the following basic items when you plan for the DS8900F power requirements:

- ▶ Power connectors
- ▶ Input voltage
- ▶ Power consumption and environment
- ▶ Power control features

Power connectors

Each DS8900F base and expansion frame features redundant intelligent Power Distribution Unit (iPDU) rack systems. The base frame can have 2 - 4 power cords, and the expansion frame two power cords.

Attach the power cords to each frame to separate AC power distribution systems. For more information about power connectors and power cords, see *IBM DS8900F Introduction and Planning Guide*, SC27-9560.

Input voltage

When you plan the power requirements of the storage system, consider the input voltage requirements. Table 5-3 and Table 5-4 shows the DS8900F input voltages and frequencies.

Table 5-3 DS8900F single-phase input voltages and frequencies

Characteristic	Voltage (single-phase)
Nominal input voltage	200 - 240 RMS V AC
Minimum input voltage	180 RMS V AC
Maximum input voltage	256 RMS V AC
Customer wall breaker rating 1-phase	30 - 63 Amps ^a
Steady-state input frequency	50 ± 3 or 60 ± 3.0 Hz
Power line disturbance (PLD) input frequencies (<10 seconds)	50 ± 3 or 60 ± 3.0 Hz

a. Can vary by region.

Table 5-4 DS8900F 3-phase input voltages and frequencies

Characteristic	Three-phase delta	Three-Phase wye
Nominal input voltage	200 - 240 RMS V AC	380 - 415 RMS V AC
Minimum input voltage	180 RMS V AC	315 RMS V AC
Maximum input voltage	256 RMS V AC	465 RMS V AC
Customer wall breaker rating 3-phase	60 - 63 Amps	20 - 32 Amps
Steady-state input frequency	50 ± 3 or 60 ± 3.0 Hz	
PLD input frequencies (<10 seconds)	50 ± 3 or 60 ± 3.0 Hz	

Power consumption

Table 5-5 lists the power consumption specifications of the DS8900F. The power estimates in this table are conservative and assume a high transaction rate workload.

Table 5-5 Power consumption and environmental information (fully equipped frames)

Measurement	Unit of measure	Base frame model	Expansion frame model
Peak electric power	Kilowatt (kW)	DS8910F model 993: 2.2 (single-phase) DS8910F model 994: 4.6 (single-phase) DS8950F model 996: 6.2 (three-phase) DS8958F model 998: 6.3 (three-phase)	Model E96: 3.9 (three-phase)
Thermal load	British thermal units (BTUs) per hour	DS8910F model 993: 7464 DS8910F model 994: 15682 DS8950F model 996: 20984 DS8980F model 998: 21489	Model E96: 13320
Capacity of exhaust	Cubic meters per min. (cubic feet per min., or CFM)	44.2 (1500)	51.8 (1800)

The values represent data that was obtained from the following configured systems:

- ▶ A standard DS8910F model 993 system that contains six sets of fully configured high-performance storage enclosures and eight Fibre Channel (FC) adapters.
- ▶ Standard base frames that contain 12 sets of fully configured high-performance storage enclosures and 16 FC adapters.
- ▶ Expansion models that contain 12 sets of fully configured high-performance storage enclosures and 16 FC adapters.

DS8900F cooling

Air circulation for the DS8900F is provided by the various fans that are installed throughout the frame. All of the fans in the DS8900F system direct air flow from the front of the frame to the rear of the frame. No air exhausts out of the top of the frame.

The use of such directional air flow allows cool aisles to the front and hot aisles to the rear of the systems, as shown in Figure 5-4.

The operating temperature for the DS8900F is 16 - 32 °C (60 - 90 °F) at relative humidity limits of 20% - 80% and optimum at 45%.

Important: The following factors must be considered when the DS8900F system is installed:

- ▶ Ensure that the air circulation for the DS8900F base frame and expansion frames is maintained free from obstruction to keep the unit operating in the specified temperature range.
- ▶ For safety reasons, do not store anything on top of the DS8900F system.




Figure 5-4 DS8900F air flow: Hot aisle/cold aisle approach

Power control features

The DS8900F has remote power control features that are used to control the power of the storage system through the HMC. For more information about power control features, see *IBM DS8900F Introduction and Planning Guide, SC27-9560*.

5.2.6 Host interface and cables

The DS8900F model 994, DS8950F model 996, and DS8980F model 998 can contain a maximum of 16 host adapters, and the DS8910F model 993 can contain a maximum of eight host adapters. The DS8950F model 996 and DS9080F model 998 allow an extra 16 host adapters to be installed in the expansion frame. For a full breakdown of the available ports on a DS8900F, see Table 2-8 on page 54.

Table 5-6 on page 151 shows the minimum and maximum numbers of host adapters that are supported by the DS8900F.

Table 5-6 Minimum and maximum host adapters

Storage system type	Storage system configuration	Minimum number of host adapter features for the base frame	Maximum number of host adapter features for the storage system
DS8910F model 993	Base	2	8
DS8910F model 994	Base frame	2	16
DS8950F model 996 and DS8980F model 998	Base frame plus expansion frames	2	32

Fibre Channel and Fibre Channel connection

Each host adapter port supports Fibre Channel Protocol (FCP) or Fibre Channel connection (IBM FICON). However, it cannot support both topologies simultaneously on the same port. Fabric components from various vendors, including IBM; Broadcom, Emulex, and Brocade; Cisco; and Marvell QLogic are supported by both environments.

The FC and FICON shortwave (SW) host adapter, when it is used with 50 μ m multi-mode fiber cable, supports point-to-point distances. For more information about cable limits, see Table 5-7.

Table 5-7 Cabling type and limits according to speed

Cable type	Distance limits		
	8-gigabit Fibre Channel (GFC)	16 GFC	32 GFC
OM2 (50 μ m)	N/A	35 m (115 ft.)	20 m (65 ft.)
OM3 (50 μ m)	150 m (492 ft.)	100 m (328 ft.)	70 m (230 ft.)
OM4 (50 μ m)	190 m (623 ft.)	125 m (410 ft.)	100 m (328 ft.)

The FC and FICON longwave (LW) host adapter, when it is used with 9 μ m single-mode fiber cable, extends the point-to-point distance to 10 km (6.2 miles).

Different fiber optic cables with various lengths can be ordered for each FC adapter port.

Table 5-8 lists the fiber optic cable features for the FCP/FICON adapters.

Table 5-8 FCP/FICON cable features

Feature Code	Length	Characteristics	Compatible FC host adapter features
1410	40 m (131 ft.)	50 μ m OM3, multimode	<ul style="list-style-type: none"> ▶ SW FC or FICON host adapters (Feature Codes 3353 and 3353) ▶ LC connector
1412	2 m (6.5 ft.)	50 μ m OM3, multimode	
1413	3 m (10 ft.)	50 μ m OM3, multimode	
1411	31 m (102 ft.)	50 μ m OM3, multimode	▶ LC/SC connector

Feature Code	Length	Characteristics	Compatible FC host adapter features
1420	31 m (102 ft.)	9 μm OS1, single mode	▶ LW FC or FICON host adapters (Feature Codes 3253 or 3257) ▶ LC connector
1422	2 m (6.5 ft.)	9 μm OS1, single mode	
1423	3 m (10 ft.)	9 μm OS1, single mode	
1421	31 m (102 ft.)	9 μm OS1, single mode	▶ LC/SC connector

For more information about IBM supported attachments, see *IBM DS8900F Introduction and Planning Guide*, SC27-9560.

For more information about host types, models, adapters, and operating systems (OSs) that are supported by the DS8900F, see the [IBM System Storage Interoperation Center \(SSIC\) for DS8000](#).

zHyperLink connections and cables

zHyperLink is a short distance IBM Z attach link that is designed for up to 10x lower latency than High-Performance FICON for IBM Z (zHPF). zHyperLink provides random reads and writes and small block sequential writes. It is a point-to-point connection that uses Peripheral Component Interconnect Express (PCIe) Gen3 with a maximum distance of 150 meters. zHyperLink connects the IBM Z (central processor complexes (CPCs)) directly to the I/O enclosure of a DS8900F system, as shown in Figure 5-5.

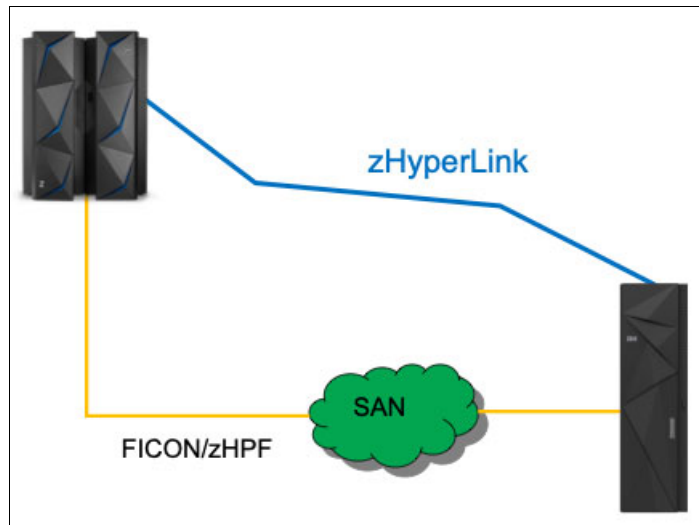


Figure 5-5 zHyperLink connection

zHyperLink does not replace zHPF. It works in cooperation with it to reduce the workload on zHPF. zHyperLink provides a new PCIe connection. The physical number of current zHPF connections is not reduced by zHyperLink.

On the DS8900F, the number of zHyperLink ports that can be installed varies, depending on the number of cores per CPC that are available and the number of I/O bay enclosures.

What happens when you enable zHyperlink on your DS8900F system?: When enabling zHyperlink write on the DS8900F for the first time, it will quiesce and resume both LPARs. Successive zHyperlink enables/disables of either write or read have no affect on DS8900F LPAR operation.

The number of zHyperLinks that can be installed based on the number of cores available is listed in Table 5-9 on page 153.

Table 5-9 zHyperLink availability by DS8900F model

System or model	Cores per CPC (DS8900F server)	zHyperLink support	Max zHyperLink connections (increments of 2)
DS8910F model 993	8	Yes	4
DS8910F model 994	8	Yes	4
DS8950F model 996	10	Yes	6
	20	Yes	8
DS8950F model 996 with expansion frame model E96	20	Yes	8 ^a
DS8980F model 998	44	Yes	8
DS8980F model 998 with expansion frame model E96	44	Yes	8 ^b

a. Maximum of 12 zHyperlink connections per system.

b. Maximum of 12 zHyperlink connections per system.

Each zHyperLink connection requires a zHyperLink I/O adapter to connect the zHyperLink cable to the storage system, as shown in Table 5-10 and Table 5-11.

Table 5-10 Feature Codes for zHyperLink I/O adapters

Feature Code	Description	Models
3500	zHyperLink I/O adapter	All

Table 5-11 shows the codes for the zHyperLink cables.

Table 5-11 Feature Codes for zHyperLink cables

Feature Code	Cable type	Cable length	Compatible zHyperLink I/O adapter features
1450	OM4 50/125 micrometer, multimode, and MTP connectors	40 m (131 Ft)	zHyperLink I/O adapter (Feature Code 3500)
1451		150 m (492 ft)	
1452		3 m (9.8 ft)	

5.2.7 Host adapter Fibre Channel specifics for open environments

Each storage system host adapter has four ports, and each port has a unique worldwide port name (WWPN). Each port can be configured as Small Computer System Interface (SCSI)-FCP or FICON topology by using the DS Management GUI or the DS CLI. Host adapters can be SW or LW. Extra host adapters up to two host adapters per I/O enclosure can be installed.

With host adapters that are configured as FC protocols, the DS8900F provides the following configuration capabilities:

- ▶ A maximum of 128 FC ports.
- ▶ A maximum of 509 logins per FC port, which includes host ports and Peer-to-Peer Remote Copy (PPRC) target and initiator ports.
- ▶ Access to 63750 logical unit numbers (LUNs) for each target (one target for each host adapter), depending on the host type.
- ▶ Either switched-fabric (FC-SW), or point-to-point topologies.
- ▶ The adapters do not support arbitrated loop topology at any speed.

5.2.8 FICON specifics on a z/OS environment

For host adapters that are configured for FICON, the DS8900F provides the following configuration capabilities:

- ▶ Fabric or point-to-point topologies
- ▶ A maximum of number of host adapter ports that depends on the model:
 - Thirty-two ports on the DS8910F model 993
 - Sixty-four ports on the DS8910F model 994, DS8950F model 996, and DS8980 model 998
 - One hundred and twenty-eight ports on the DS8950F model 996 and DS8980F with model E96 expansion frame
- ▶ A maximum of 509 logins for each host adapter port
- ▶ A maximum of 8,192 logins for each storage unit
- ▶ A maximum of 1,280 logical paths on each host adapter port
- ▶ Access to all 255 control-unit images (65,280 Count Key Data (CKD) devices) over each FICON port
- ▶ A maximum of 512 logical paths for each control-unit image

Note: The IBM z16™, z15, z14, and z13 servers support 32,000 devices for each FICON host channel. The IBM zEnterprise® EC12 and IBM zEnterprise BC12 servers support 24,000 devices for each FICON host channel. Earlier IBM Z servers support 16,384 devices for each 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 an FC switch or FICON director to a single storage system.

5.2.9 Best practices for host adapters

For optimum availability and performance, the following practices are preferred:

- ▶ To obtain the maximum ratio for availability and performance, install one host adapter on each available I/O enclosure before you install the second host adapter on the same I/O enclosure.
- ▶ The DS8900F supports 16 or 32 GFC 4-port host adapters. Based on the configuration, these host adapters or an intermix of them can be installed in the DS8900F.

Note: IBM z16 now supports 32 GFC host adapters and FICON Express32S Channels, which provide twice the read/write bandwidth compared to 16 GFC host adapters on previous models, thus taking full advantage of 32 GFC host adapters on the DS8900F.

- ▶ Better performance for copy services can be obtained by using dedicated host ports for remote copy links, and other path optimization. For more information, see *IBM DS8900F Performance Best Practices and Monitoring*, SG24-8501.

Note: DS8000 has a set of internal parameters that are known as *pokeables*, which sometimes are referred to as product switches. These internal parameters are set to provide the best behavior in most typical environments. In special cases, like intercontinental distances or when bandwidth is low, some internal tuning might be required to adjust those internal controls to keep Global Mirror (GM) as efficient as it is in more common environments. Pokeable values can be displayed by a GUI or by Copy Services Manager, but they can be changed only by IBM Support. For more information, see *DS8000 Global Mirror Best Practices*, REDP-5246.

5.2.10 Worldwide node name and worldwide port name determination

The incoming and outgoing data to the DS8900F is tracked by using a worldwide node name (WWNN) and a WWPN. For the DS8000, each storage facility image (SFI) has its own unique WWNN. The storage unit itself also has a unique WWNN. Each host adapter port has a unique and persistent WWPN for attachment to a SAN. The WWNN and WWPN values can be determined by using the DS CLI or DS Storage Management GUI.

Determining a WWNN by using a DS CLI

The DS8900F WWNN has an address that is similar to the following strings:

```
50:05:07:63:0z:FF:Cx:xx  
50:50:07:63:0z:FF:Dx:xx
```

The z and x:xx values are unique combinations for each system and each SFI that are based on a machine's serial number. Use the DS CLI command `lssi` to determine the SFI WWNN, as shown in Example 5-1.

Example 5-1 SFI WWNN determination

```
dsccli> lssi  
Date/Time: June 23, 2022 5:37:34 PM CEST IBM DSCLI Version: 7.9.30.154 DS: -  
Name      ID              Storage Unit      Model WWNN              State  ESSNet  
=====
```

ds8k-r9-11	IBM.2107-75LLB71	IBM.2107-75LLB70	998	500507630AFFD3E7	Online	Enabled
------------	------------------	------------------	-----	------------------	--------	---------

Do not use the `lssu` command because it determines the storage unit WWNN, which is not used. Attached hosts see only the SFI, as shown in Example 5-2.

Example 5-2 Machine WWNN

```
dsccli> lssu  
Date/Time: June 23, 2022 5:42:31 PM CEST IBM DSCLI Version: 7.9.30.154 DS: -  
Name ID              Model WWNN              pw state  
=====
```

-	IBM.2107-75LLB70	998	500507630AFFEBE7	On
---	------------------	-----	------------------	----

Determining a WWPN by using a DS CLI

Similar to the WWNN, a WWPN in the DS8900F looks like the following address:

50:05:07:63:0z:YY:Yx:xx

However, the DS8900F WWPN is a child of the SFI WWNN, where the WWPN inserts the z and x:xx values from SFI WWNN. It also includes the YY:Y from the logical port naming, which is derived from where the host adapter is physically installed. Use the DS CLI command `lσιοport` to determine the SFI WWPN, as shown in Example 5-3.

Example 5-3 WWPN determination

```
dscli> lσιοport
Date/Time: June 23, 2022 5:43:43 PM CEST IBM DSCLI Version: 7.9.30.154 DS: IBM.2107-75LLB71
ID      WWPN              State              Type              topo      portgrp Security
=====
I0200 500507630A1013E7 Communication established Fibre Channel-SW SCSI-FCP 0      Disabled
I0201 500507630A1053E7 Communication established Fibre Channel-SW FICON    0      Disabled
I0202 500507630A1093E7 Communication established Fibre Channel-SW FICON    0      Disabled
I0203 500507630A10D3E7 Communication established Fibre Channel-SW FICON    0      Disabled
I0230 500507630A1313E7 Offline              Fibre Channel-LW -      0      Disabled
I0231 500507630A1353E7 Offline              Fibre Channel-LW -      0      Disabled
I0232 500507630A1393E7 Offline              Fibre Channel-LW -      0      Disabled
I0233 500507630A13D3E7 Offline              Fibre Channel-LW -      0      Disabled
I0240 500507630A1413E7 Communication established Fibre Channel-SW FICON    0      Disabled
I0241 500507630A1453E7 Offline              Fibre Channel-SW -      0      Disabled
I0242 500507630A1493E7 Communication established Fibre Channel-SW SCSI-FCP 0      Disabled
I0243 500507630A14D3E7 Offline              Fibre Channel-SW -      0      Disabled
I0300 500507630A1813E7 Offline              Fibre Channel-LW -      0      Disabled
I0301 500507630A1853E7 Offline              Fibre Channel-LW -      0      Disabled
I0302 500507630A1893E7 Offline              Fibre Channel-LW -      0      Disabled
I0303 500507630A18D3E7 Offline              Fibre Channel-LW -      0      Disabled
I0310 500507630A1913E7 Communication established Fibre Channel-SW SCSI-FCP 0      Disabled
I0311 500507630A1953E7 Communication established Fibre Channel-SW FICON    0      Disabled
I0312 500507630A1993E7 Communication established Fibre Channel-SW FICON    0      Disabled
I0313 500507630A19D3E7 Communication established Fibre Channel-SW FICON    0      Disabled
I0330 500507630A1B13E7 Communication established Fibre Channel-SW FICON    0      Disabled
I0331 500507630A1B53E7 Offline              Fibre Channel-SW -      0      Disabled
I0332 500507630A1B93E7 Communication established Fibre Channel-SW SCSI-FCP 0      Disabled
I0333 500507630A1BD3E7 Offline              Fibre Channel-SW -      0      Disabled
```

Determining a WWNN by using a DS GUI

Use the following steps to determine the WWNN by using the Storage Management GUI:

1. Connect by using a web browser to the HMC IP address:

`https://< hmc IP address >`

2. Select **Actions**.
3. Select **Properties** to obtain the WWNN value, as shown in Figure 5-6 on page 157.




Figure 5-6 System properties: WWNN

You can also determine the host adapter port WWPN by completing the following steps:

1. Connect to the HMC IP address by using a web browser:
`https://< hmc IP address >`
2. Select **Actions**.
3. Select **Modify Fibre Channel Port Protocols**.

The default view may show protocols and the state only. The view can be customized to display the port WWPN and the frame.

- Click **Actions**, and then select **Customize Columns** to include the WWPN and frame in the view. You receive the full list of each installed I/O port with its WWPN and its physical location, as shown in Figure 5-7.

ID	Host Adapter	State	Protocol
I0200		✓ Communication established	SCSI FCP
I0201		✓ Communication established	FICON
I0202		✓ Communication established	FICON
I0203		✓ Communication established	FICON
I0230		✓ Unconfigured	Undefined
I0231		✓ Unconfigured	Undefined
I0232		✓ Unconfigured	Undefined
I0233		✓ Unconfigured	Undefined
I0240		✓ Communication established	FICON
I0241		✓ Unconfigured	Undefined
I0242		✓ Communication established	SCSI FCP
I0243		✓ Unconfigured	Undefined
I0300		✓ Unconfigured	Undefined
I0301		✓ Unconfigured	Undefined
I0302	1	✓ Unconfigured	Undefined

Figure 5-7 Determining the I/O port WWPN

You can also select **Show System Health Overview** and then **Fibre Channel Ports**, as shown in Figure 5-8.

ID	Frame	State	Protocol	WWPN
I0203	1	✓ Communication established	FICON	50:05:07:6:
I0202	1	✓ Communication established	FICON	50:05:07:6:
I0201	1	✓ Communication established	FICON	50:05:07:6:
I0200	1	✓ Communication established	SCSI FCP	50:05:07:6:
I0240	1	✓ Communication established	FICON	50:05:07:6:
I0242	1	✓ Communication established	SCSI FCP	50:05:07:6:
I0312	1	✓ Communication established	FICON	50:05:07:6:
I0310	1	✓ Communication established	SCSI FCP	50:05:07:6:
I0313	1	✓ Communication established	FICON	50:05:07:6:
I0311	1	✓ Communication established	FICON	50:05:07:6:
I0332	1	✓ Communication established	SCSI FCP	50:05:07:6:

Figure 5-8 System Health Overview: Fibre Channel Ports

5.3 Network connectivity planning

To implement the DS8900F, you must consider the physical network connectivity of the HMC within your LAN.

Consider the following network and communications requirements when you plan the location and interoperability of your storage systems:

- ▶ HMC network access (one IP per HMC).
- ▶ Remote support connection.
- ▶ SAN connectivity.
- ▶ An IBM Security Guardium Key Lifecycle Manager connection if encryption, end-point security, or TCT is activated, or an LDAP connection if LDAP is implemented.

For more information about physical network connectivity, see *IBM DS8900F Introduction and Planning Guide*, SC27-9560.

5.3.1 Hardware Management Console and network access

The HMCs are the focal point for the configuration of Advanced Function management, and maintenance for a DS8900F unit. Two internal HMCs are included with every base frame.

A dual Ethernet connection is available for client access. The two HMCs provide redundant management access to enable continuous availability access for encryption key servers and other advanced functions.

The HMC can be connected to the client network for the following tasks:

- ▶ Remote management of your system by using the DS CLI
- ▶ Remote management by using the DS Storage Management GUI by opening a browser to the network address of the HMC:

```
https://<HMC IP address>
```

To access the HMCs (HMC1/HMC2) over the network, provide the following information:

- ▶ HMC: For each HMC, determine one TCP/IP address, hostname, and domain name.
- ▶ DNS settings: If a DNS is implemented, ensure that it is reachable to avoid contention or network timeout.
- ▶ Gateway routing information: Supply the necessary routing information.

Note: Users also can control a second Ethernet adapter within the HMCs. This capability is available only by Request for Price Quotation (RPQ).

For more information about HMC planning, see Chapter 6, “IBM DS8900F Management Console planning and setup” on page 167.

Important: The DS8900F uses 172.16.y.z and 172.17.y.z private network addresses. If the client network uses the same addresses, the IBM SSR can reconfigure the private networks to use another address range option.

5.3.2 IBM Spectrum Control and IBM Storage Insights

IBM Spectrum Control is an integrated software solution that can help you improve and centralize the management of storage environments. With IBM Spectrum Control, you can manage and configure multiple DS8000 storage systems from a single point of control.

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

IBM Storage Insights is a complementary offering to IBM Spectrum Control.

IBM Storage Insights is offered free of charge to customers who own IBM block storage systems. It is an IBM Cloud storage service that monitors IBM block storage. It provides single-pane views of IBM block storage systems, such as the Operations dashboard and the Notifications dashboard.

With the information that is provided, such as the diagnostic event information; key capacity; and performance information, and the streamlined support experience, you can quickly assess the health of your storage environment and get help with resolving issues.

On the Advisor window, IBM Storage Insights provides recommendations about the remedial steps that can be taken to manage risks and resolve issues that might impact your storage services. For a brief illustration of IBM Storage Insights features, see 12.11, “Using IBM Storage Insights” on page 448.

5.3.3 DS Command-Line Interface

You can use the DS CLI can be used to create, delete, modify, and view CS functions, and to perform the logical configuration of a storage unit. These tasks can be performed interactively, in batch processes (OS shell scripts), or by using DS CLI script files.

A *DS CLI script file* is a text file that contains one or more DS CLI commands. It can be issued as a single command. DS CLI can also be used to manage other functions for a storage system, including managing security settings, querying point-in-time performance information or the status of physical resources, and exporting audit logs.

The DS CLI client can be installed on a workstation, and can support multiple OSs. The DS CLI client can access the DS8900F over the client’s network. For more information about hardware and software requirements for the DS CLI, see *IBM DS8000 Series Command-Line Interface User’s Guide*, SC27-9562.

5.3.4 Remote support connection

Remote support is available through the embedded AOS application or through the IBM Remote Support Center (RSC).

Embedded AOS

The preferred remote support connectivity method for IBM is through Transport Layer Security (TLS) for the Management Console (MC) to IBM communication. DS8900F uses an embedded AOS server solution. Embedded AOS is a secure and fast broadband form of remote access.

IBM Remote Support Center

DS8900F also supports a web-based remote support option that is called RSC. It provides more isolation between IBM Support and the DS8900F system by requiring that all remote support actions be performed by using a web-based console interface.

For more information, see Chapter 6, “IBM DS8900F Management Console planning and setup” on page 167 and Chapter 12, “Monitoring and support” on page 423.

5.3.5 Storage area network connection

The DS8900F can be attached to a SAN environment through its host adapter ports. The SAN provides the capability to interconnect open systems hosts, IBM Z hosts, and other storage systems.

A SAN allows your host bus adapter (HBA) host ports to have physical access to multiple host adapter ports on the storage system. Zoning can be implemented to limit the access (and provide access security) of host ports to the storage system.

Shared access to a storage system host adapter port is possible from hosts that support a combination of HBA types and OSs.

Important: A SAN administrator must verify periodically that the SAN is working correctly before any new devices are installed. SAN bandwidth must also be evaluated to ensure that it can handle the new workload.

5.3.6 Key manager servers for encryption

The DS8900F system is delivered with FDE drives. When you activate encryption, either isolated key managers or local key managers are required. The Local Key Management feature is available with DS8900 Release 9.2 or later, and new systems can select the local key manager feature (Feature Code #0405) at the time of ordering. It is a chargeable feature.

With a DS8900F system, you can choose among IBM Security Guardium Key Lifecycle Manager, Gemalto SafeNet KeySecure, and Thales Vormetric Data Security Manager for data at rest and TCT encryption. IBM Fibre Channel Endpoint Security encryption requires IBM Security Guardium Key Lifecycle Manager. You cannot mix IBM Security Guardium Key Lifecycle Manager and SafeNet or Vormetric key servers. For more information, see *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security (DS8000 Release 9.2)*, REDP-4500.

Encryption activation review planning

IBM Encryption offerings must be activated before they are used. This activation is part of the installation and configuration steps that are required to use the technology.

Using the IBM Security Guardium Key Lifecycle Manager

An IBM Security Guardium Key Lifecycle Manager license is required to use the IBM Security Guardium Key Lifecycle Manager software. Two isolated Guardium Key Lifecycle Manager servers are required to enable data at rest encryption on the DS8900F system.

Important: Clients must acquire an IBM Security Guardium Key Lifecycle Manager license to use the Guardium Key Lifecycle Manager software.

Note: The licensing for IBM Security Guardium Key Lifecycle Manager includes both an installation license for the Guardium Key Lifecycle Manager management software and licensing for the encrypting drives.

The DS8000 series supports IBM Security Guardium Key Lifecycle Manager V2.5 or later. This version also uses a connection between the HMC and the key server, which complies with the National Institute of Standards and Technology (NIST) SP 800-131A standard. For TCT encryption, IBM Security Guardium Key Lifecycle Manager V3.0.0.2 or later is required. For IBM Fibre Channel Endpoint Security encryption, IBM Security Guardium Key Lifecycle Manager V4.0 or later is required.

You are advised to upgrade to the latest version of the IBM Security Guardium Key Lifecycle Manager.

IBM Security Guardium Key Lifecycle Manager connectivity and routing information

To connect the IBM Security Guardium Key Lifecycle Manager to your network, provide the following settings to your IBM SSR:

- ▶ IBM Security Guardium Key Lifecycle Manager server network IDs, hostnames, and domain name
- ▶ DNS settings (if you plan to use DNS to resolve network names)

Two network ports must be opened on a firewall to allow the DS8900F connection and to obtain an administration management interface to the IBM Security Guardium Key Lifecycle Manager server. These ports are defined by the IBM Security Guardium Key Lifecycle Manager administrator.

For more information, see the following IBM publications for IBM Security Guardium Key Lifecycle Manager:

- ▶ *IBM Security Guardium Key Lifecycle Manager Quick Start Guide*, GI13-4178
- ▶ *IBM Security Key Lifecycle Manager Installation and Configuration Guide*, SC27-5335, or the relevant sections in [IBM Security Guardium Key Lifecycle Manager 4.1.0](#).

Local Key Manager

With Release 9.2 or later, it is possible to configure Local Key Management. For more information about Local Key Management, see *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security (DS8000 Release 9.2)*, REDP-4500.

5.3.7 Lightweight Directory Access Protocol server

The DS8000 system provides, by default, local basic user authentication. The user IDs, roles, and their passwords are maintained locally within the DS8000 system. Each individual DS8000 system has its own list of user IDs and passwords that must be maintained separately.

To simplify user ID management and comply with industry or company-internal security regulations, the DS8000 system can access a centralized directory service to perform user authentication by using LDAP.

Since Release 9.1, LDAP authentication can be configured through the Storage Management GUI, as described in 6.5.2, “Remote authentication” on page 188.

The native LDAP implementation does not require the IBM Copy Services Manager proxy. For more information, see *LDAP Authentication for IBM Storage DS8000 Systems: Updated for DS8000 Release 9.3.2*, REDP-5460.

5.4 Remote Mirror and Remote Copy connectivity

The DS8900F uses the high-speed FCP for Remote Mirror and Remote Copy connectivity. Ensure that you assigned sufficient FCP paths for the remote mirroring between the source and target sites to address performance and redundancy issues. When you plan to use Metro Mirror (MM) and Global Copy (GC) modes between a pair of storage systems, use separate logical and physical paths for the MM, and use another set of logical and physical paths for the GC.

Plan the distance between the primary and auxiliary storage systems carefully to correctly acquire fiber optic cables of the necessary length that are required. If necessary, the CS solution can include hardware, such as channel extenders or dense wavelength division multiplexing (DWDM).

For more information, see *IBM DS8000 Copy Services: Updated for IBM DS8000 Release 9.1*, SG24-8367.

5.5 Disk capacity considerations

The effective capacity of the DS8900F is determined by the following factors:

- ▶ The spare configuration
- ▶ The capacity of the installed drives
- ▶ The selected RAID configuration: RAID 6, RAID 10, or when applicable RAID 5
- ▶ The storage type: Fixed-Block (FB) or CKD

5.5.1 Disk sparing

RAID arrays automatically attempt to recover from a drive failure by rebuilding the data for the failed drive to a spare disk drive module (DDM). For sparing to occur, a drive with a capacity equal to or greater than the failed drive must be available on the same device adapter (DA) pair. After the sparing is initiated, the spare and the failing drives are swapped between their respective array sites so that the spare drive becomes part of the array site that is associated with the array at the failed drive. The failing drive becomes a failed spare drive in the array site from which the spare came.

High-Performance Flash Enclosure Gen2

High-Performance Flash Enclosures (HPFEs) are installed in pairs, with 16, 32, or 48 flash drives per enclosure pair. Two spare flash drives are assigned for each HPFE Gen2 pair. If a flash drive fails and a spare is taken, the system calls for service because only one spare remains in the HPFE Gen2 pair (DA pair).

For more information about the DS8000 sparing concepts, see 3.5.11, “Spare creation” on page 96.

5.5.2 Disk capacity

The following RAID configurations are supported on the DS8900F:

- ▶ 5+P+Q+S RAID 6 configuration: The array consists of five data drives and two parity drives. The remaining drive on the array site is used as a spare.
- ▶ 6+P+Q RAID 6 configuration: The array consists of six data drives and two parity drives.
- ▶ 3+3+2S RAID 10 configuration: The array consists of three data drives that are mirrored to three copy drives. Two drives on the array site are used as spares.
- ▶ 4+4 RAID 10 configuration: The array consists of four data drives that are mirrored to four copy drives.
- ▶ 6+P+S RAID 5 configuration (by RPQ only): The array consists of six data drives and one parity drive. The remaining drive of the array site is used as a spare.
- ▶ 7+P RAID 5 configuration (by RPQ only): The array consists of seven data drives and one parity drive.

Note: The following characteristics refer to RAID:

- ▶ Spare drives are globally available to the flash RAID controller pair.
- ▶ The +P/+Q indicators do not mean that a single drive is dedicated to holding the parity bits for the RAID. The DS8900F uses floating parity technology so that no single drive is always involved in every write operation. The data and parity stripes float among the member drives of the array to provide optimum write performance.

For the effective capacity of one rank in the various possible configurations, see *IBM DS8900F Introduction and Planning Guide*, SC27-9560.

Important: When you review the effective capacity, consider the following points:

- ▶ Effective capacities are in decimal gigabytes (GB). 1 GB is 1,000,000,000 bytes.
- ▶ Although drive sets contain 16 drives, arrays use only eight drives. The effective capacity assumes that you have two arrays for each disk drive set.

The IBM Storage Modeller tool can help you determine the raw and net storage capacities and the numbers for the required extents for each available type of RAID. IBM Storage Modeller is available only for IBM employees and IBM Business Partners.

DS8900F offers the following flash drives sets with HPFE Gen2:

- ▶ 2.5-inch high-performance flash (Tier 0) drives are 800 TB, 1.6 TB, and 3.2 TB capacity drives.
- ▶ 2.5-inch high-capacity flash (Tier 1) drives are 3.84 TB.
- ▶ 2.5-inch high-capacity flash (Tier 2) drives are 1.92 TB, 7.68 TB, and 15.36 TB capacity drives.

Flash drives in HPFE Gen2 are ordered in sets of 16 within an enclosure pair. There are three sets of 16 drives in an HPFE Gen2 enclosure pair.

Important: The following restrictions apply:

- ▶ An RPQ or Storage Customer Opportunity Request (SCORE) is required to use RAID 5.
- ▶ Within one HPFE Gen2 pair of six array sites, a RAID intermix is allowed, but no intermix of high-performance drives (Flash Tier 0) with high-capacity drives (Flash Tier 1 or Flash Tier 2) is supported.

For the latest information about supported RAID configurations and requesting an RPQ or SCORE, contact your IBM SSR.



IBM DS8900F Management Console planning and setup

This chapter describes the planning tasks that are involved in the setup of the required IBM DS8900F Management Console (MC), which is also known as the Hardware Management Console (HMC).

This chapter covers the following topics:

- ▶ DS8900F Management Console overview
- ▶ Management Console software
- ▶ Management Console activities
- ▶ Management Console network settings
- ▶ User management
- ▶ Secondary Management Console

6.1 DS8900F Management Console overview

The MC is a multi-purpose piece of equipment that provides the services to configure and manage storage, and manage several of the operational aspects of the storage system. It also provides the interface where service personnel perform diagnostic and repair tasks.

The MC does not process any of the data from hosts. It is not even in the path that the data takes from a host to the storage. The MC is a configuration and management station for the whole DS8900F system.

The DS8900F includes a Management Enclosure (ME). This enclosure contains two MCs, which is standard for redundancy reasons, but the ME contains other essential management components too, which are explained in 6.1.1, “Management Enclosure” on page 168.

The MC, which is the focal point for DS8900F management, includes the following functions:

- ▶ DS8900F power control
- ▶ Storage provisioning
- ▶ Storage system health monitoring
- ▶ Storage system performance monitoring
- ▶ Copy Services (CS) monitoring
- ▶ Embedded IBM Copy Services Manager
- ▶ Interface for onsite service personnel
- ▶ Collection of diagnostic and Call Home data
- ▶ Problem management and alerting
- ▶ Enables remote support access
- ▶ Storage management through the DS GUI
- ▶ Connection to IBM Security Guardium Key Lifecycle Manager or other supported external key manager for encryption management functions, if required
- ▶ Connection to an external IBM Copy Services Manager or IBM Spectrum Control
- ▶ Interface for Licensed Internal Code (LIC) and other firmware updates

6.1.1 Management Enclosure

The ME is a 2U chassis containing the following components:

- ▶ Two MCs (MCs or HMCs) as standard
- ▶ Two Ethernet switches
- ▶ Two rack power control cards (RPCCs)
- ▶ Two power supply units (PSUs) to power the ME components
- ▶ One Local or Remote switch assembly
- ▶ Internal cabling for communications and power for each of the components

Figure 6-1 shows a ME.




Figure 6-1 Management Enclosure

Note: The location of the ME can be slightly different from what is shown in Figure 6-1 because there are many rack configurations, such as IBM DS8980F model 998, IBM DS8950F model 996, IBM DS8910F model 994, and IBM DS8910F Rack-Mounted model 993 that can fit into an existing 19-inch form-factor rack. On racked DS8900F systems, the ME is always in the base frame (Rack 1).

The ME is designed to create a compact container for all essential system management components that otherwise would be mounted around the rack as in former IBM DS8000 models.

Figure 6-2 shows the layout of the components of the ME.




Figure 6-2 Management Enclosure component layout

The ME provides internal communications to all of the modules of the DS8900F system. It also provides external connectivity by using two Ethernet cables from each HMC for remote management, and provides keyboard, mouse, and video connectivity from each HMC for local management. Cables are routed from the MCs to the rear of the ME through a cable management arm (CMA).

6.1.2 Management Console hardware

The MC itself consists of a small form-factor (SFF) personal computer running Red Hat Linux. Using an SFF personal computer makes the MC efficient in many ways, including power consumption. It supports the DS8900F hardware and firmware installation and maintenance activities.

Because of the small width, both primary and secondary MCs are mounted in the front of the ME next to each other.

There is an 1U keyboard and display tray that are available. For racked DS8980F, DS8950, and DS8910F model 994 systems, you must order one (use Feature Code 1765). For the Flexibility Class Rack-Mounted model 993, it is optional. For more information, see *IBM DS8910F Model 993 Rack-Mounted Storage System Release 9.1*, REDP-5566.

The MC connects to the customer network and provides access to functions that can be used to manage the DS8900F. Management functions include logical configuration, problem notification, Call Home for service, remote service, and CS management.

These management functions can be performed from the DS GUI, Data Storage Command-Line Interface (DS CLI), or other storage management software that supports the DS8900F.

For example, clients who use an external IBM Copy Services Manager for advanced functions, such as Metro Mirror (MM) or FlashCopy are communicating to the storage system by connecting the IBM Copy Services Manager server to the HMCs as management entry point.

The MC provides connectivity between the DS8000 and Encryption Key Manager (EKM) servers (Security Guardium Key Lifecycle Manager), and also provides the functions for remote call home and remote support connectivity.

The MCs are equipped with Ethernet connections for the client's network. For more information, see 6.1.3, "Private and Management Ethernet networks" on page 171.

To provide continuous availability to the MC functions, the DS8900F includes a second MC by default. The secondary HMC is needed for redundancy, such as for encryption management or CS functions. For more information about the secondary MC, see 6.6, "Secondary Management Console" on page 197.

HMC hardware revision

The amount of memory, the solid-state drive (SSD) capacity, and the type of processor that is used in the HMCs might change with the DS8900F release.

Use the DS CLI command `lshmc` to show the HMC types, whether both HMCs are online, and their amount of disk capacity and memory, as shown in Example 6-1.

Example 6-1 The lshmc command

```
dsccli> lshmc
```

Name	State	Role	Release	Management-IP	Location-Code	Machine-Type	Machine-Model	Memory	Disk
d1r20	Online	Primary(1)	R9.2 bundle	89.20.131.0 10.11.236.76	U1700L47.I355527	1700	L47	32 GB	256 GB
d1r20-b	Online	Secondary(2)	R9.2 bundle	89.20.131.0 10.11.236.77	U1700L47.I355529	1700	L47	32 GB	256 GB

6.1.3 Private and Management Ethernet networks

The HMCs communicate with the storage facility internally through a pair of redundant Ethernet networks, which are designated as the *black network* and the *gray network*. There are two switches that are isolated from each other, and they are in the ME.

Inside the ME, the switch ports of the internal black and gray network switches are routed from inside the ME to an external breakout at the rear of the ME by using short patch cables to make the ports accessible from outside.

Each central processor complex (CPC) flexible service processor (FSP) and each CPC logical partition (LPAR) network are connected to both switches. Each of these components (FSP and LPAR) uses their own designated interface for the black network and another interface for the gray network. These components are connected to the external breakout ports of the ME.

Each MC also uses two designated Ethernet interfaces for the internal black (eth0) and gray (eth3) networks. The MCs that are installed in the ME are already connected internally to the switches without routing them to connections outside of the ME.

Additionally, an MC contains a third Ethernet interface (eth2) for the customer network connection to allow management functions to be started over the network. This customer network connection is routed from the MC directly to the rear of ME to its own breakout ports.

For particular circumstances where the customer needs a second Ethernet interface for management reasons, you can place a Request for Price Quotation (RPQ) to have a USB Ethernet adapter (eth1) added to the MC. This adapter can be used to connect the HMC to two separate customer networks, usually for separating internet traffic (call home and remote access) from storage management tasks (DS GUI, DS CLI, and IBM Copy Services Manager).

Figure 6-3 shows these internal and external network connections.

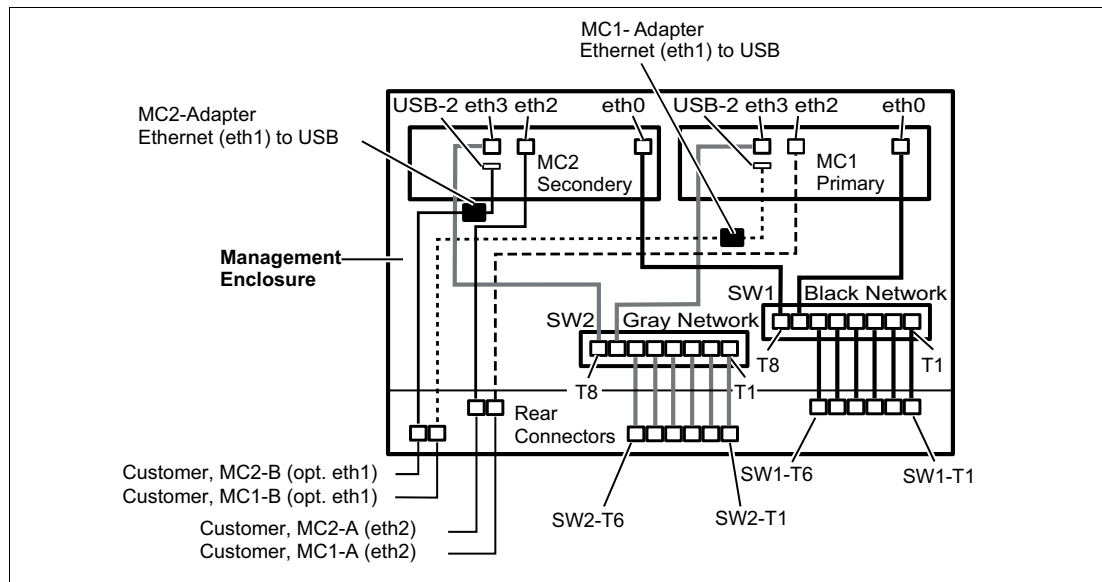


Figure 6-3 ME and MC internal network connections

Important: The internal Ethernet switches that are shown in Figure 6-3 and Figure 6-4 on page 173 (the black and gray private networks) are for DS8900F internal communication only. Do not connect these ports directly to your network. There is no connection between the customer network interfaces and the black and gray network to keep them isolated.




Figure 6-4 ME external connections

Intelligent Power Distribution Units

In racked configurations such as DS8980F, DS8950F, and DS8910F model 994 systems, the rack power control is managed through the Ethernet-managed intelligent Power Distribution Units (iPDUs). For rackless configurations such as DS8910F model 993, the iPDUs are optional.

An HMC communicates to these iPDUs by using the Ethernet network, and it manages and monitors the system power state, iPDU configuration, System AC power on and off, iPDU firmware update, iPDU health check and error, and power usage reporting.

The iPDUs' network interfaces are also connected to the external ports of the ME to reach the black and gray network switches. They are distributed over the black and gray network, which means iPDUs that belong to one power domain (usually on the left side of the rear of the rack) connect to a gray network switch and the iPDUs that belong to the other power domain (usually on the right side of the rear of the rack).

For the rack-mounted DS8980F and DS8950F systems, you may add an expansion frame model E96. Two cascaded switches are added to the base frame to connect the additional iPDUs of the expansion rack to the ME switches.

One 1U 24-port Ethernet switch is added for the black network and one 24-port is added for the gray network. Each of them has an uplink to the related ME switch port of their designated black or gray network. The 2U space that is required is already reserved at the bottom of the base rack.

Figure 6-5 and Table 6-1 show how each port is used on the pair of DS8900F ME Ethernet switches.

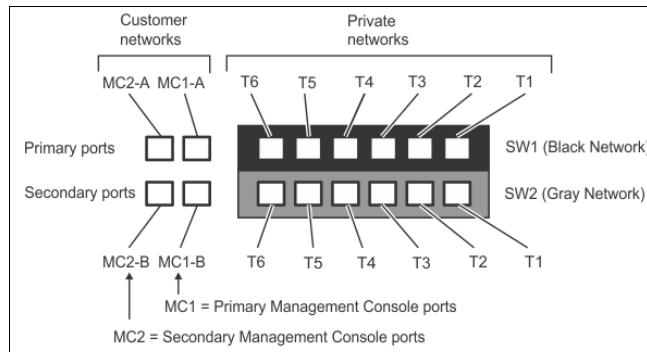


Figure 6-5 ME rear Ethernet breakout

Table 6-1 Functional connections of the switch ports

From ME rear outside connectors	Connects to
SW1-T1 black	Unused or iPDU-E22 (upper right from rear) or uplink to 24-port cascaded Ethernet switch black network
SW2-T1 gray	Unused or iPDU-E21 (upper left from rear) or uplink to 24-port cascaded Ethernet switch gray network
SW1-T2 black	Unused or iPDU-E24 lower right from rear
SW2-T2 gray	Unused or iPDU-E23 lower left from rear
SW1-T3 black / SW2-T3 gray	Upper CPC/CEC1 FSP black / gray
SW1-T4 black / SW2-T4 gray	Lower CPC/CEC1 FSP black / gray
SW1-T5 black / SW2-T5 gray	Upper CPC/CEC1 LPAR black / gray
SW1-T6 black / SW2-T6 gray	Lower CPC/CEC2 LPAR black / gray
SW1-T7 black / SW2-T7 gray	MC1 eth0 black / MC1 eth3 gray
SW1-T8 black / SW2-T8 gray	MC2 eth0 black / MC2 eth3 gray

6.2 Management Console software

The MC, which is based on Linux, includes the following application servers:

- ▶ DS8000 Storage Management GUI
- ▶ National Institute of Standards and Technology (NIST) Web UI (WUI)
- ▶ IBM Copy Services Manager
- ▶ RESTful application programming interface (API) services
- ▶ DS CLI
- ▶ IBM Enterprise Storage Server® Network Interface (IBM ESSNI) server

The Management Console also provides the interfaces for IBM Spectrum Control, IBM Storage Insights, and the DS CLI to connect to the DS8900F remotely.

Note: The DS Open API with IBM System Storage Common Information Model (CIM) agent is no longer supported. The removal of the CIM Agent simplifies network security because fewer open ports are required.

6.2.1 DS Storage Management GUI

The DS GUI is used to perform logical configuration and storage management tasks. It can be accessed in the following ways:

- ▶ Remotely by using a web browser that is connected to the DS8900F HMC (or MC)
- ▶ From the local console of the HMC directly
- ▶ By using IBM Spectrum Control, which has connectivity to the HMC

For more information, see 9.2.1, “Accessing the DS GUI” on page 235.

6.2.2 Data Storage Command-Line Interface

The Data Storage Command-Line Interface (DS CLI), which must be run in the command environment of an external workstation, is a second option to communicate with the MC. The DS CLI is a good choice for configuration tasks when many updates are needed. A copy of DS CLI is installed locally on the HMC, and it can be used when servicing the machine from the local console.

Note: The DS Storage Management GUI also provides a built-in DS CLI. Look for the console icon on the lower left of the browser window after logging in.

For more information about DS CLI usage and configuration, see Chapter 10, “IBM DS8900F Storage Management Command-line Interface” on page 339. For a complete list of DS CLI commands, see *IBM DS8000 Series: Command-Line Interface User’s Guide*, SC27-9562.

6.2.3 RESTful application programming interface

DS8900F RESTful API services provide an easy-to-use application programming interface (API) to manage DS8900F through communication with the MC. The RESTful API communicates with RESTful services that run on the MC. The RESTful services in turn interact with the IBM ESSNI server software that runs on the MC to pass requests and receive replies. For more information about the RESTful API, see *IBM DS8880/DS8870 RESTful API Guide*, SC27-9235.

6.2.4 IBM Copy Services Manager interface

IBM Copy Services Manager is preinstalled on the DS8900F MC. You can use it to manage and automate replication and disaster recovery (DR) for up to four DS8000 storage systems.

This feature removes the requirement for an external server to host IBM Copy Services Manager, which provides savings on infrastructure costs and operating system (OS) licensing. Administration costs are also reduced because the embedded IBM Copy Services Manager instance is upgraded through the DS8900F code maintenance schedule, which is performed by IBM support personnel.

Important: Avoid configuring the primary HMC and the secondary HMC of the same storage system as the active and standby IBM Copy Services Manager servers within a CS environment.

6.2.5 Updating the embedded IBM Copy Services Manager

With the embedded IBM Copy Services Manager on HMC, the IBM Copy Services Manager release that came initially installed on the DS8000 HMC might be outdated. IBM Copy Services Manager can be updated independent of a microcode update.

Important: Updating the HMC embedded IBM Copy Services Manager must be done exclusively through the IBM DS CLI tool that is installed on the workstation, laptop, or server.

Update IBM Copy Services Manager on the HMC by completing the following steps:

1. Verify the current level of the DS CLI.
2. Verify the current level of IBM Copy Services Manager on the HMC.
3. Download selected releases of DS CLI, if necessary, and IBM Copy Services Manager from [IBM Fix Central](#).
4. Update DS CLI, if needed.
5. Update IBM Copy Services Manager on the HMC.

The [DS8000 Code Recommendation](#) page provides a link to the DS8900F code bundle information page, as shown in Figure 6-6 and Figure 6-7.

Model	Minimum Installed Level	Recommended "go to" Level	Latest Level
DS8900F (R9.2)	R9.2 GA - 89.20.147.0 [Aug 2021] (45, 47, 48)	R9.2 SP1.2 - 89.21.31.0 [Jan 2022] (47, 48)	R9.2 SP1.5 - 89.21.35.0 [Apr 2022]

Figure 6-6 DS8900F Code Recommendation page

DS8900F Bundle	SEA or LMC Version	DSCLI Client	Heat Map Transfer Utility	Storage Manager	Copy Services Manager
89.21.35.0	7.9.21.98	7.9.21.80	7.9.21.80	5.9.21.1035	6.3.2

Figure 6-7 Recommended DS CLI release

Verifying the current level of IBM Copy Services Manager on the HMC

To verify the current IBM Copy Services Manager release that is installed on a DS8000 HMC, run the `lsoftware` DS CLI command:

```
lsoftware -l -type csm -hmc all
```

Example 6-2 on page 177 shows an example where the IBM Copy Services Manager release on both HMCs is 6.2.9.1.

Example 6-2 Current IBM Copy Services Manager release

```
dsccli> lssoftware -l -type csm -hmc all
Type Version                Status HMC
=====
CSM V6.2.9.1-a20200804-1704 Running 2
CSM V6.2.9.1-a20200804-1704 Running 1
dsccli>
```

Downloading IBM Copy Services Manager for an upgrade on the HMC

The IBM Copy Services Manager installation file must be downloaded on the same workstation or server where the DS CLI was previously installed.

Complete the following steps. Assume that IBM Copy Services Manager 6.3.0 is the release that will be installed.

1. On the [IBM Fix Central](#) page, select **IBM Copy Services Manager** as the product, **6.3.0.0** as the installed version, and **Linux** as the platform. Figure 6-8 shows a summary of selected options.

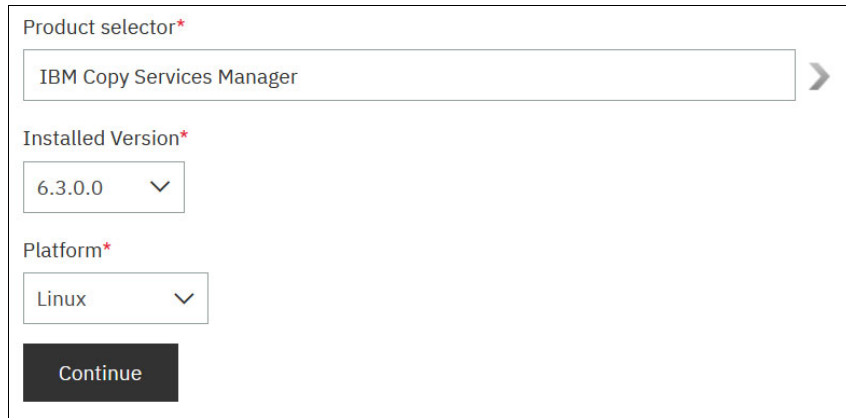


Figure 6-8 Selected IBM Copy Services Manager Version for HMC

Note: The HMC OS is Linux. Ensure that the correct platform is selected.

2. Be sure to download the correct Linux-x86_64 release. Figure 6-9 shows the correct package type selected. Check the Release Notes, and if there is a newer fix pack file, you can use it instead.

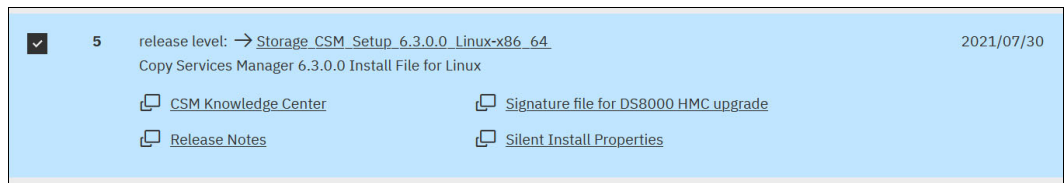


Figure 6-9 IBM Copy Services Manager Linux-x86_64 package

- When the download process is complete, note the folder path where the files were stored. In our example, the files are stored in the folder C:\Downloads\CSM_Linux, as shown in Figure 6-10.

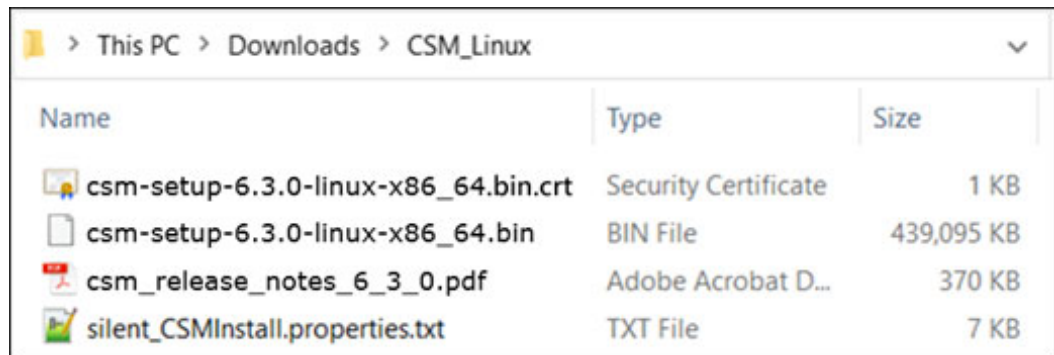


Figure 6-10 Downloaded IBM Copy Services Manager files

Updating IBM Copy Services Manager on the HMC by using the DS CLI

Update the IBM Copy Services Manager on each HMC. In a dual HMC environment, update one IBM Copy Services Manager instance at a time.

Note: If your IBM Copy Services Manager installation has active CS sessions, you must follow [best practices](#) while applying maintenance to an active management server.

Note: The Active and Standby servers must be updated concurrently. Failure to do so results in the inability to connect to the other server.

The DS CLI command that is used for the IBM Copy Services Manager update is `installsoftware`. You can find more information about the command in [IBM Documentation](#).

Table 6-2 describes the parameters that are necessary for the `installsoftware` command.

Table 6-2 DS CLI `installsoftware` parameters

Parameter	Explanation
<code>-type csm</code>	The software type of the installation package.
<code>-loc software_package</code>	The full path of the installation package to be installed.
<code>-certloc certificate_location</code>	The full path of the certificate file location.
<code>-hmc 1 2 all</code>	Specifies the primary or secondary HMC where the software is to be installed. The default is <code>all</code> .

Note: Ensure that no spaces are included in the path that you specify for the location of the software package and certificate file.

Note: In addition to the standard 1751 port, DS CLI also uses port 1755 (TCP protocol) to transfer the IBM Copy Services Manager installation file to the HMC. That port must be open on any physical or software firewall standing between the workstation where DS CLI is installed and the DS8000 HMCs.

To effectively run the command, you must use a DS8000 user ID that is part of the Administrator role (for example, the default admin user ID).

Example 6-3 shows how the IBM Copy Services Manager on HMC1 was updated by using DS CLI.

Example 6-3 IBM Copy Services Manager update on HMC1

```
dscli> installsoftware -type csm -loc
C:\Downloads\CSM_Linux\csm-setup-6.3.0-linux-x86_64.bin -certloc
C:\Downloads\CSM_Linux\csm-setup-6.3.0-linux-x86_64.bin.crt -hmc 1
CMUC00516I installsoftware: The file uploaded successfully.
CMUC00517I installsoftware: Software CSM is successfully installed on 1.
```

```
dscli> lssoftware
Type Version                Status
=====
CSM V6.3.0.0-a20210622-1237 Running
CSM V6.2.9.1-a20200804-1704 Running
```

The next step is to update IBM Copy Services Manager on HMC2, as shown in Example 6-4.

Example 6-4 IBM Copy Services Manager update on HMC2

```
dscli> installsoftware -type csm -loc
C:\Downloads\CSM_Linux\csm-setup-6.3.0-linux-x86_64.bin -certloc
C:\Downloads\CSM_Linux\csm-setup-6.3.0-linux-x86_64.bin.crt -hmc 2
CMUC00516I installsoftware: The file uploaded successfully.
CMUC00517I installsoftware: Software CSM is successfully installed on 2.
```

```
dscli> lssoftware
Type Version                Status
=====
CSM V6.3.0.0-a20210622-1237 Running
CSM V6.3.0.0-a20210622-1237 Running
```

The IBM Copy Services Manager upgrade is now complete.

6.2.6 Web User Interface

The HMC Web User Interface (WUI) is used for the initial setup of the HMC and for servicing the hardware of the DS8900F. It provides remote access to system utilities.

To log in to the WUI, complete the following steps:

1. Start the Storage Management GUI, as shown in Figure 6-11. Click the **Service** icon (wrench) to access the Service MC.

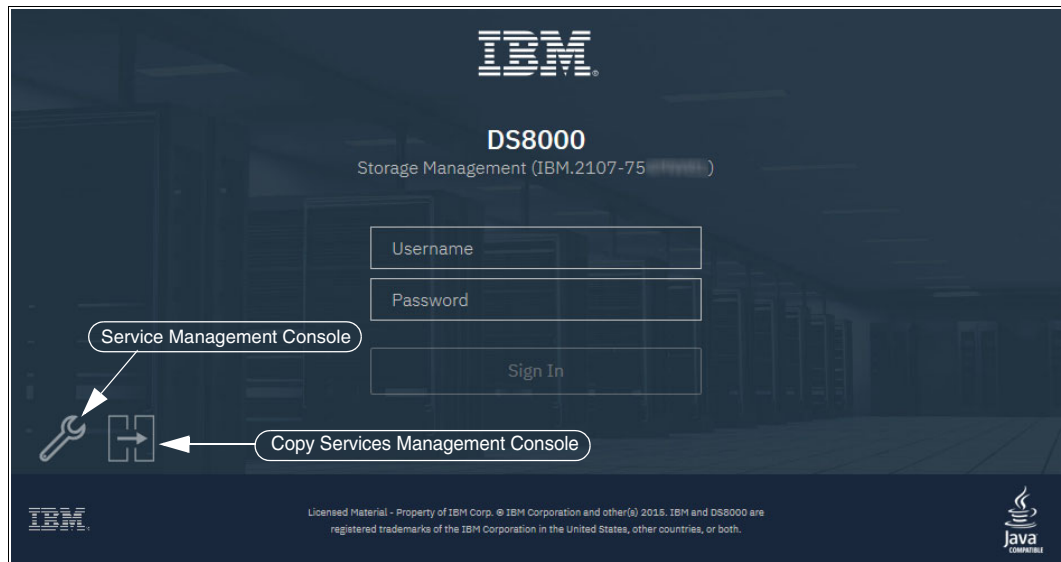


Figure 6-11 DS Storage Management GUI Logon window

2. Click **Log on and launch the Hardware Management Console web application** to open the login window, as shown in Figure 6-12 on page 181, and log in. The default user ID is customer and the default password is cust0mer.

Important: Make sure to change the default password. The user credentials for accessing the Service Management Console (HMC) are managed separately from the ones that are used with DS CLI and the Storage Management GUI. For more information about HMC user management, see 6.5.3, “Service Management Console User Management” on page 189.

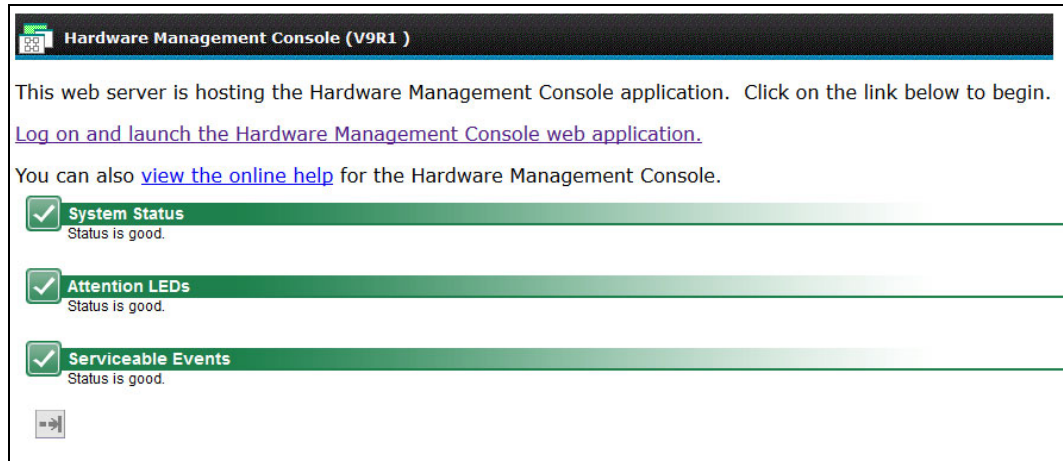


Figure 6-12 Service Management Console application

3. If you are successfully logged in, you see the MC window, in which you can select **Status Overview** to see the status of the DS8900F. Other areas of interest are shown in Figure 6-13.

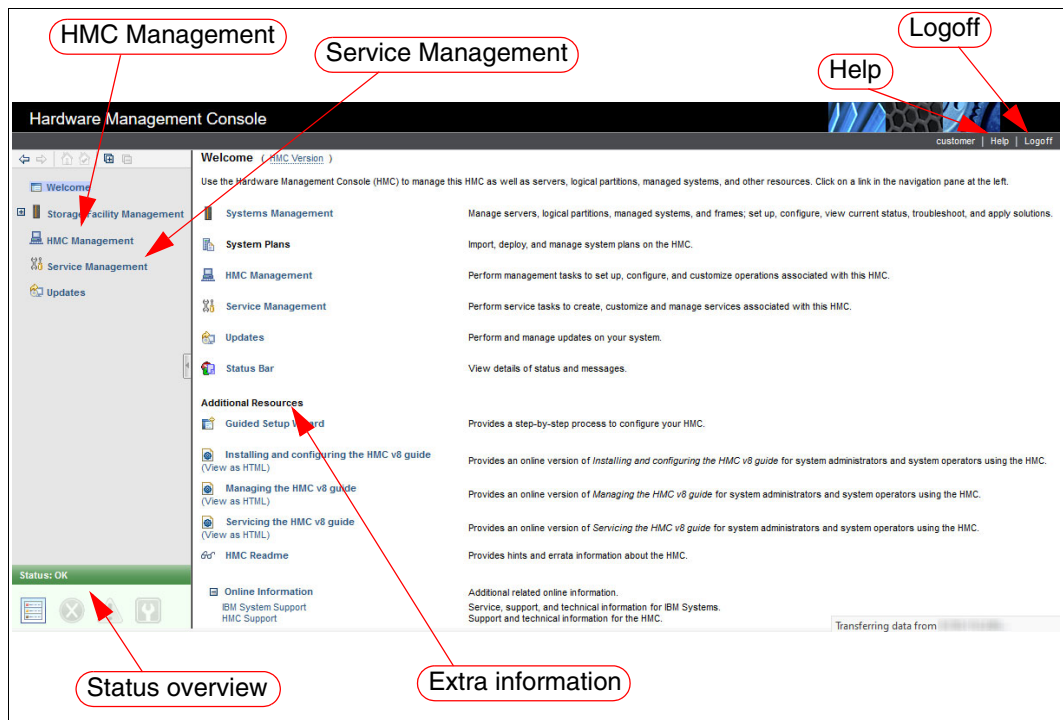


Figure 6-13 Web UI main window

Because the MC web UI is mainly a services interface, it is not covered here. For more information, see the **Help** menu.

6.2.7 IBM ESSNI server

IBM ESSNI is the logical server that communicates with the DS GUI server and interacts with the two processor nodes of the DS8900F. It is also referred to as the DS Network Interface (DSNI).

6.3 Management Console activities

This section covers planning and maintenance tasks for the DS8900F MC. For more information about overall planning, see Chapter 5, “IBM DS8900F physical planning and installation” on page 141.

6.3.1 Management Console planning tasks

To plan the installation or configuration of the MC, you must do the following tasks:

- ▶ A connection to the client network is needed at the base rack for the primary MC. Another connection is also needed at the location of the secondary MC. The connections must be standard CAT5/6 Ethernet cabling with RJ45 connectors.
- ▶ IP addresses for the primary and secondary MCs are needed. The DS8900F can work with IPv4 and IPv6 networks. For more information about procedures to configure the DS8900F MC for IPv6, see 6.4, “Management Console network settings” on page 185.
- ▶ Most users access the DS GUI remotely through a browser. You can also use IBM Spectrum Control in your environment to access the DS GUI.
- ▶ The web browser to be used on any administration workstation must be supported, as described in *IBM DS8900F Introduction and Planning Guide, SC27-9560*.
- ▶ The IP addresses of Simple Network Management Protocol (SNMP) recipients must be identified if the client wants the DS8900F MC to send traps to a monitoring station.
- ▶ Email accounts must be identified if the client wants the DS8900F MC to send email messages for problem conditions.
- ▶ The IP addresses of Network Time Protocol (NTP) servers must be identified if the client wants the DS8900F MC to use NTP for time synchronization.
- ▶ When a DS8900F is ordered, the license and certain optional features must be activated as part of the customization of the DS8900F. For more information, see Chapter 7, “IBM DS8900F features and licensed functions” on page 199.
- ▶ The installation tasks for the optional external MC must be identified as part of the overall project plan and agreed upon with the responsible IBM personnel.

Important: Applying feature activation codes is a concurrent action.

6.3.2 Planning for Licensed Internal Code upgrades

The following tasks must be considered regarding the LIC upgrades on the DS8900F:

- ▶ LIC changes

IBM periodically releases changes to the DS8900F series Licensed Machine Code (LMC). Customers can check the IBM Support site for the latest [Flashes, Alerts and Bulletins](#), and keep up to date by subscribing to [IBM Support Notifications](#).

- ▶ LIC installation options

There are three installation types available, depending on the support contract:

- On-site Code Load

An IBM Systems Service Representative (IBM SSR) goes onsite to install the changes.

- Remote Code Load (RCL)

IBM Remote Support personnel install the LIC remotely.

- Customer Code Load

As of Release 9.3, customers can perform the installation.

- ▶ DS CLI Compatibility

Check whether the new LIC requires new levels of DS CLI. Plan on upgrading them on the relevant workstations, if necessary.

- ▶ Code prerequisites

When you are planning for initial installation or for LIC updates, ensure that all prerequisites for the environment are identified correctly, which include host OS versions, fixes, host bus adapter (HBA) levels, interconnect and fabric types, and OS versions.

DS8900F interoperability information is available at the [IBM System Storage Interoperation Center \(SSIC\)](#).

To prepare for downloading the drivers, see the “Interoperability Search Details” report in SSIC, which provides an end-to-end support matrix from the host to the DS8900F, and covers all versions of OS, multipathing software, and firmware. This check is necessary to ensure that the DS8900F storage subsystem is in a supported environment.

Important: The SSIC includes information about the latest supported code levels. This availability does not necessarily mean that former levels of HBA firmware or drivers are no longer supported. Some host type interoperability, such as NetApp ONTAP, might need to be confirmed in the vendor’s support matrix. If you are in doubt about any supported levels, contact your IBM SSR.

Never proceed with a LIC update without adhering to all prerequisites.

- ▶ Maintenance windows

The LIC update of the DS8900F is a nondisruptive action. Scheduling a maintenance window with added time for contingency is still a best practice. Also, plan for sufficient time to confirm that all environment prerequisites are met before the upgrade begins.

For more information about LIC upgrades, see Chapter 11, “Licensed Machine Code” on page 405.

6.3.3 Time synchronization

With the DS8900F, the MC can use the NTP service. Clients can specify NTP servers on their internal or external network to provide the time to the MC. It is the client's responsibility to ensure that the NTP servers are working, stable, and accurate. An IBM SSR enables the MC to use NTP servers (ideally at the time of the initial DS8900F installation). Changes can be made by the client by using the **Change Date and Time** action under **MC Management** on the MC.

Important: For correct error analysis, the date and time information must be synchronized on all components in the DS8900F environment. These components include the DS8900F MC, the attached hosts, IBM Spectrum Control, and DS CLI workstations.

6.3.4 Monitoring DS8900F with the Management Console

A client can receive notifications from the MC through traps and email messages. Notifications contain information about your storage complex, such as open serviceable events. You can choose one or both of the following notification methods:

► Traps

For monitoring purposes, the DS8900F uses traps. A trap can be sent to a server in the client's environment, perhaps with System Management Software, which handles the trap that is based on the Management Information Base (MIB) that was delivered with the DS8900F software. A MIB that contains all of the traps can be used for integration purposes into System Management Software.

The supported traps are described in the documentation that comes with the LIC on the CDs that are provided by the IBM SSR. The IP address to which the traps must be sent must be configured during initial installation of the DS8900F. For more information about the DS8900F and monitoring, see Chapter 12, "Monitoring and support" on page 423.

► Email

When you enable email notifications, email messages are sent to all the addresses that are defined on the MC whenever the storage complex encounters a serviceable event or must alert individuals to other information.

During the planning process, create a list of the individuals who need to be notified.

Additionally, when the DS8900F is attached to an IBM Z system server, a service information message (SIM) notification occurs automatically. A SIM message is displayed on the OS console if a serviceable event occurs. These messages are not sent from the MC, but from the DS8900F through the channel connections that run between the server and the DS8900F.

6.3.5 Event notification through syslog

To meet ever increasing security requirements, the DS8900F supports security and logging events that are forwarded to a syslog server. This capability was previously available only on the MC. Events that are contained in the audit log are forwarded to configured syslog receivers.

Up to eight external syslog servers can be configured, with varying ports if required. Events that are forwarded include user login and logout, all commands that are issued by using the GUI or DS CLI while the user is logged in, and remote access events. Events are sent from Facility 19, and are logged as level 6.

6.3.6 Call Home and remote support

The MC uses outbound (Call Home) and inbound (remote service) support.

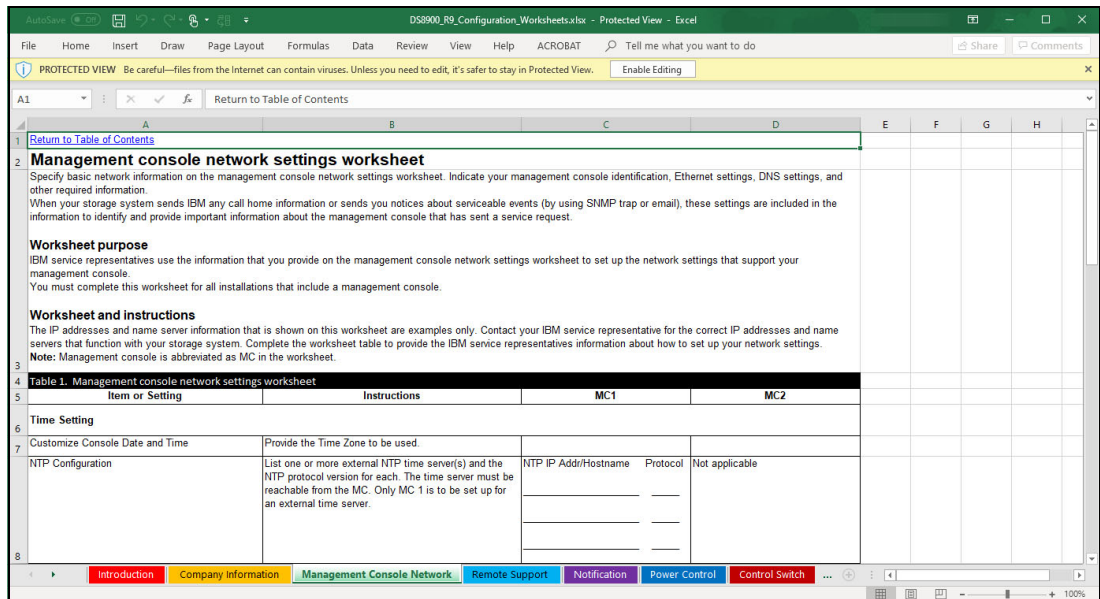
Call Home is the capability of the MC to contact the IBM Support Center to report a serviceable event. Remote support is the capability of IBM SSR to connect to the MC to perform service tasks remotely. If the IBM Support Center can connect to the MC to perform service tasks remotely based on the setup of the client's environment, an IBM SSR can connect to the MC to perform detailed problem analysis. The IBM SSR can view error logs and problem logs and start trace or memory dump retrievals.

Remote support can be configured by using the embedded Assist On-site (AOS) or Remote Support Console. The setup of the remote support environment is performed by the IBM SSR during the initial installation. For more information, see Chapter 12, "Monitoring and support" on page 423.

6.4 Management Console network settings

The DS8900F MC is configured by an IBM SSR during the initial installation of the storage facility. The IBM SSR applies the settings that are defined by the customer in the [DS8000 Configuration Worksheets](#), which are partially shown in Figure 6-14.

This activity includes the configuration of the private (internal) and management (customer) network with IPv6 or IPv4, hostname, DNS, NTP, routing, and remote support settings.



The screenshot shows an Excel spreadsheet with the following content:

Return to Table of Contents

Management console network settings worksheet

Specify basic network information on the management console network settings worksheet. Indicate your management console identification, Ethernet settings, DNS settings, and other required information.

When your storage system sends IBM any call home information or sends you notices about serviceable events (by using SNMP trap or email), these settings are included in the information to identify and provide important information about the management console that has sent a service request.

Worksheet purpose

IBM service representatives use the information that you provide on the management console network settings worksheet to set up the network settings that support your management console.

You must complete this worksheet for all installations that include a management console.

Worksheet and instructions

The IP addresses and name server information that is shown on this worksheet are examples only. Contact your IBM service representative for the correct IP addresses and name servers that function with your storage system. Complete the worksheet table to provide the IBM service representatives information about how to set up your network settings.

Note: Management console is abbreviated as MC in the worksheet.

Table 1. Management console network settings worksheet

Item or Setting	Instructions	MC1	MC2
Time Setting			
Customize Console Date and Time	Provide the Time Zone to be used.		
NTP Configuration	List one or more external NTP time server(s) and the NTP protocol version for each. The time server must be reachable from the MC. Only MC 1 is to be set up for an external time server.	NTP IP Addr/Hostname Protocol _____ _____ _____	Not applicable

Figure 6-14 DS8000 Configuration Worksheets

Chapter 8, "Configuration flow" on page 225 explains the configuration flow in more detail.

Those settings can be changed afterward by using the Service Management Console WUI or DS GUI.

Configuring the Management Console Network

In the Service Console (WUI), select **HMC Management** → **Change Network Settings** → **LAN Adapters**, select the adapter, and then select **Details**, as shown in Figure 6-15.

Note: Only the customer management network interfaces eth2 and eth1 are shown and can be configured in the Network Settings dialog because the internal private black and gray networks with interfaces eth0 and eth3 are used for the running system. The eth0 and eth3 interfaces can be changed only by opening an IBM support request.

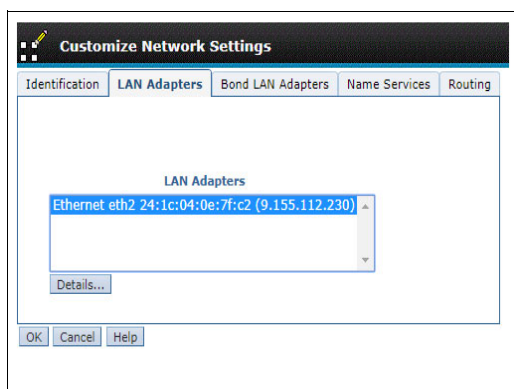


Figure 6-15 Management Ethernet settings

6.4.1 Private networks

The internal private networks (black and gray) are using the IP ranges 172.x.x.x by default. You must ensure that the network that is used for the internal private network is not interfering with outside network ranges that are used by any network that the MC can reach.

If the default address range cannot be used because it conflicts with another network, you can instead specify one of three optional addresses ranges. Table 6-3 shows the possible options that can be chosen during installation.

Table 6-3 Private networks

Setting	Black network (MC eth0)	Gray network (MC eth3)
Default	172.16.0.0 - 172.16.255.255	172.17.0.0 - 172.17.255.255
Option 1	10.235.0.0 - 10.235.2.255	10.236.0.0 - 10.236.2.255
Option 2	192.168.160.0 - 192.168.162.255	192.168.240.0 - 192.168.242.255
Option 3	9.15.0.0 - 9.15.2.255	9.16.0.0 - 9.16.2.255

When you change the internal private network, you do not need to configure each individual network interface. Instead, each change that you make changes both the black and gray networks at once.

To make the change, select **HMC Management** → **Query/Change IP Range**, as shown in Figure 6-16.

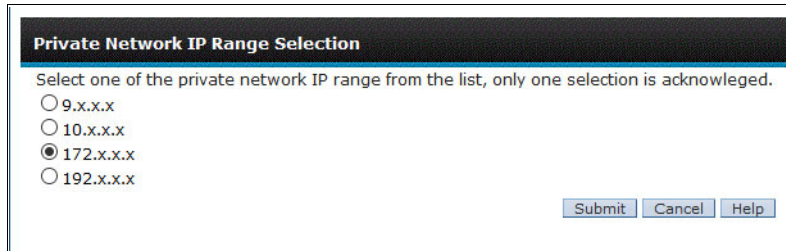


Figure 6-16 Setting the private network range

Note: Changing the internal private network range on the storage system facility can be done in concurrent mode, but requires special care. For that reason, an IBM service request must be opened before making such a change.

6.5 User management

The Service Management GUI uses credentials that are separate from the Storage Management GUI and DS CLI. This section describes user management for the DS GUI and DS CLI first. User management of the Service Management Console is described in 6.5.3, “Service Management Console User Management” on page 189.

To manage the DS GUI and DS CLI credentials, you can use the DS CLI or the DS GUI. An administrator user ID is preconfigured during the installation of the DS8900F and this user ID uses the following defaults:

- ▶ User ID: admin
- ▶ Password: admin

The password of the admin user ID must be changed before it can be used. The GUI forces you to change the password when you first log in. By using the DS CLI, you log in but you cannot run any other commands until you change the password. For example, to change the admin user’s password to passw0rd, run the following DS CLI command:

```
chuser -pw passw0rd admin
```

After you issue that command, you can run other commands.

6.5.1 Password policies

DS8900F supports different role-based users. For more information about user and role management, see 8.2, “User and role management” on page 226. When the administrator adds a user, the administrator enters a password. During the user’s first login, this password must be changed. Password settings include the period (in days) after which passwords expire and a number that identifies how many failed logins are allowed. The user ID is deactivated if an invalid password is entered more times than the limit. Only a user with administrator rights can then reset the user ID with a new initial password.

Recommendation: Do not set the value of the **chpass** command to 0 because this setting indicates that passwords never expire and unlimited login attempts are allowed.

If access is denied for the admin user, for example, because of the number of invalid login attempts, the administrator can use the security recovery utility tool on the MC to reset the password to the default value. The detailed procedure is described by selecting **Help Contents** and can be accessed from the DS GUI.

Important: Upgrading an existing storage system to the latest code release does not change the old default user-acquired rules. Existing default values are retained to prevent disruption. The user might opt to use the new defaults by running the **chpass -reset** command. The command resets all default values to the new defaults immediately.

The password for each user account is forced to adhere to the following rules:

- ▶ Passwords must contain one character from at least two groups of the following ones: alphabetic, numeric, and punctuation.
- ▶ The range for minimum password length is 6 - 64 characters. The default minimum password length is 8 characters.
- ▶ Passwords cannot contain the user's ID.
- ▶ Passwords are case-sensitive.
- ▶ The length of the password is determined by the administrator.
- ▶ Initial passwords on new user accounts are expired.
- ▶ Passwords that are reset by an administrator are expired.
- ▶ Users must change expired passwords at the next logon.

The following password security implementations are included:

- ▶ Password rules are checked when passwords are changed.
- ▶ The valid character set, embedded user ID, age, length, and history are also checked.
- ▶ Passwords that are invalidated by a change remain usable until the next password change.
- ▶ Users with invalidated passwords are not automatically disconnected from the DS8900F.
- ▶ The following password rules are checked when a user logs on:
 - Password expiration, locked-out user, and failed attempts are checked.
 - Users with passwords that expire or that are locked out by the administrator while they are logged on are not automatically disconnected from the DS8900F.

6.5.2 Remote authentication

You can enable and configure remote authentication to connect to an LDAP repository.

Starting with Release 9.1 the remote authentication setup can be found in the Storage Manager GUI. Go to the **Access** menu and select **Remote Authentication**. From there, click **Configure Remote Authentication**. The installation is guided by the Remote Authentication wizard.

DS8900F now has native support for Remote Authentication through LDAP, although it is still supported to use IBM Copy Services Manager servers as a proxy to the remote authentication servers.

Figure 6-17 shows the window that opens directly after the Welcome window. After you complete all the wizard steps of the wizard, the DS8000 is enabled and configured for remote authentication.

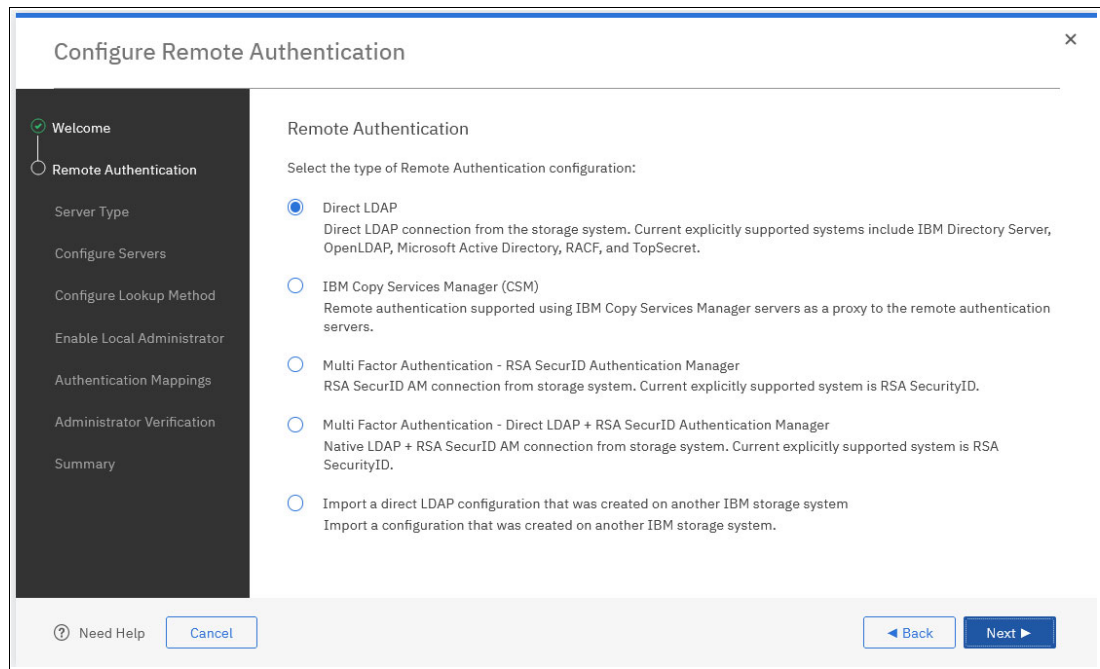


Figure 6-17 Remote Authentication wizard

The following prerequisites are required to complete the Remote Authentication wizard:

- ▶ Access to create users and groups on your remote authentication server.
- ▶ A primary LDAP repository URI is required.
- ▶ A secondary LDAP repository URI is optional.
- ▶ A User search base (only for Direct LDAP).
- ▶ A truststore file with a password is required (only for IBM Copy Services Manager).
- ▶ An IBM WebSphere® username with a password is required (only for IBM Copy Services Manager).

For more information about LDAP-based authentication and configuration, see *LDAP Authentication for IBM Storage DS8000 Systems: Updated for DS8000 Release 9.3.2*, REDP-5460.

6.5.3 Service Management Console User Management

Access to the Service Management Console is managed through the HMC WUI. With the HMC, you can manage users, user roles, and authentication methods. Creating personal IDs enables individual accountability. The HMC also supports remote authentication and centralized user ID and password control through LDAP.

In the HMC Management section of the WUI, two options are available:

- ▶ Managed User Profiles and Access
- ▶ Configure LDAP

Figure 6-18 shows the two options.

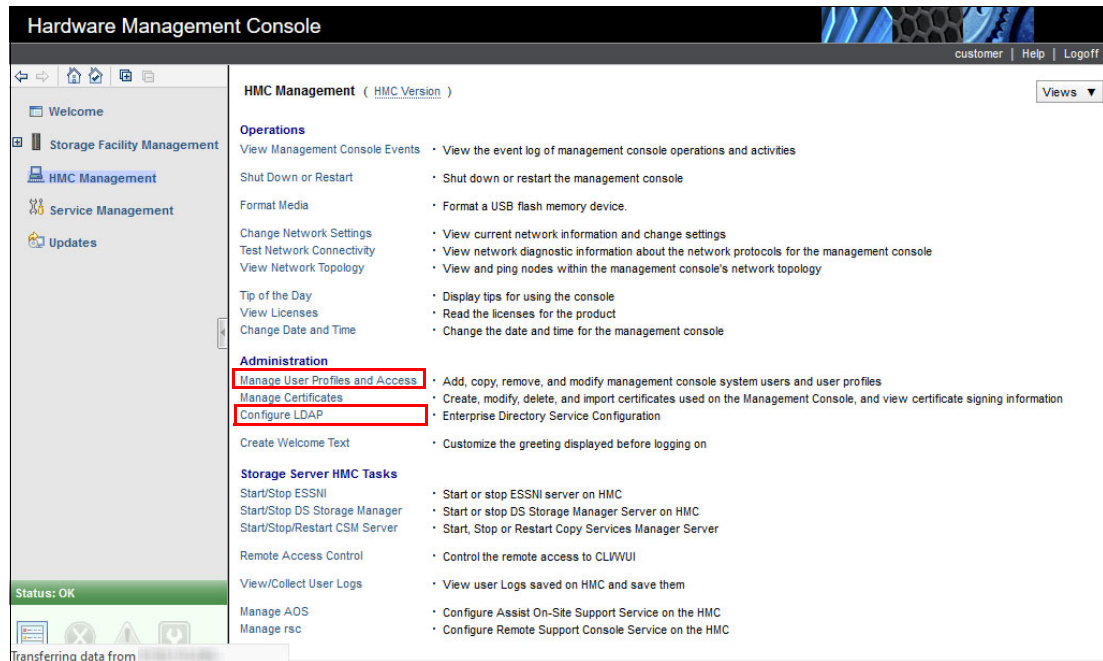


Figure 6-18 Administration options

When logged in as an admin user, you can:

- ▶ Create user IDs for predefined roles, including Service and Engineering roles.
- ▶ Modify or remove any user, including predefined Service and Engineering user IDs.
- ▶ Allow each user to change their password.
- ▶ Allow the Service and Engineering user IDs to connect remotely.
- ▶ Configure authentication locally or by using LDAP.

Important: Do not delete the last user ID in a role. For more information about removing user IDs for a role, see Table 6-5 on page 191.

There are three predefined user roles that are related to the Customer, Service, and Engineering user IDs, as shown in Table 6-4.

Table 6-4 Predefined user roles

Predefined user role	Access requirement
esshmccustomer	Requires a password for access regardless of authentication method.
esshmcserv	Local access only. Requires an IBM Support Representative to be at the HMC.
esshmcpce	Requires the IBM proprietary challenge/response key for remote access.

The roles, access, and properties for each user ID are described in Table 6-5 on page 191.

Table 6-5 User roles

Role	esshmccustomer	esshmcserv	esshmcpce
Access	Administration	Service (CE / IBM SSR)	Service (IBM Remote Support Center (RSC) and IBM SSR)
Default user ID	customer	IBM use only	IBM use only
Default Password	cust0mer	IBM use only	IBM use only
Remove last user in this role	No	Yes ^a	Yes ^b
Backup and restore in the event of HMC rebuild	Yes	Yes	Yes
LDAP Authentication	Yes	Yes ^c	Yes ^c

- a. If removed, IBM service personnel cannot perform service functions.
- b. Removing this user ID prevents remote access for support services. This user ID should not be modified or deleted.
- c. This user ID can log in to this account only by using the IBM proprietary challenge/response process. Extra user IDs with this role do *not* use the challenge/response process and are not viable for support services.

Manage User Profile and Access windows

To manage user profiles, complete the following steps:

1. Log in to the web UI, as explained in 6.2.6, “Web User Interface” on page 180.
2. In the HMC Management window, select **Manage User Profiles and Access**, as shown in Figure 6-19.

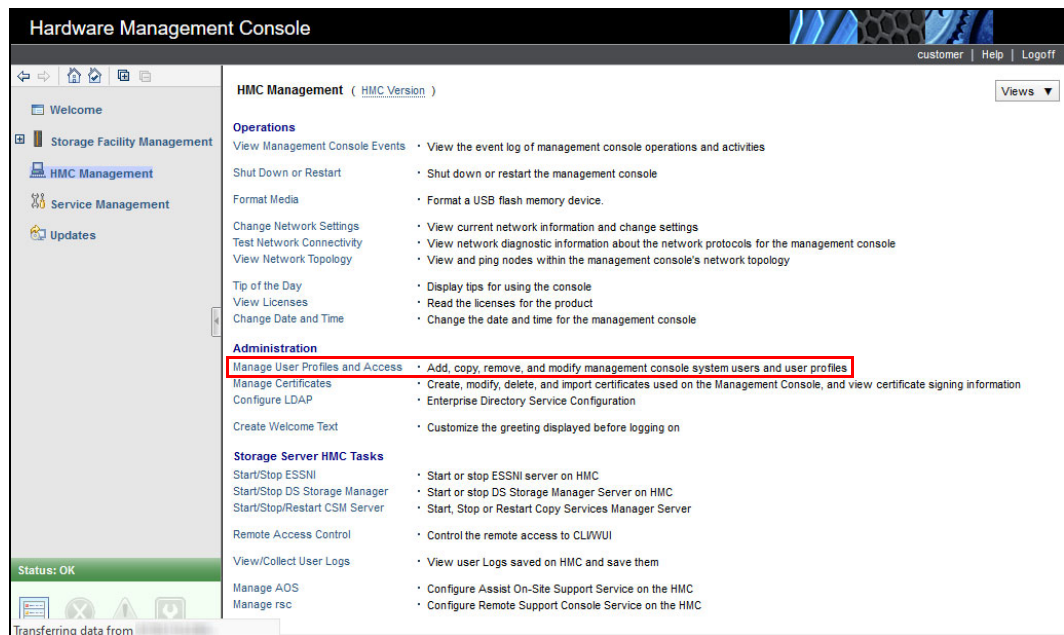


Figure 6-19 Manage User Profiles and Access

A new window opens that lists the user IDs and profiles for the defined console users, as shown in Figure 6-20.

Select	User ID	Task Role	Resource Role	Authentication Type	Description
<input checked="" type="radio"/>	hscroot	hmcsuperadmin	AllSystemResources	Local Authentication	HMC Super User
<input type="radio"/>	CE	esshmcserv	AllSystemResources	Local Authentication	HMC User
<input type="radio"/>	PE	esshmcpce	AllSystemResources	Local Authentication	HMC User
<input type="radio"/>	customer	esshmccustomer	AllSystemResources	Local Authentication	HMC User
<input type="radio"/>	essbase	hmcsuperadmin	AllSystemResources	Local Authentication	HMC User
<input type="radio"/>	root	hmcsuperadmin	AllSystemResources	Local Authentication	root

Figure 6-20 Predefined User Profiles

User IDs PE, CE, and customer are specifically for DS8900F use. Ignore the other profiles.

Note: Do not change the user ID PE because it uses the remote challenge/response login process, which is logged and audited.

The user ID root cannot log in to the WUI. The user IDs hscroot and essbase cannot access HMC functions externally. Do not use them.

Do not create user IDs with a Task Role beginning with “hmc”.

Do not create user IDs with a Task Role of esshmcooperator.

Adding a user ID

To add a user ID, complete the following steps:

1. Click **User** to open the User Profiles option menu, as shown in Figure 6-21.

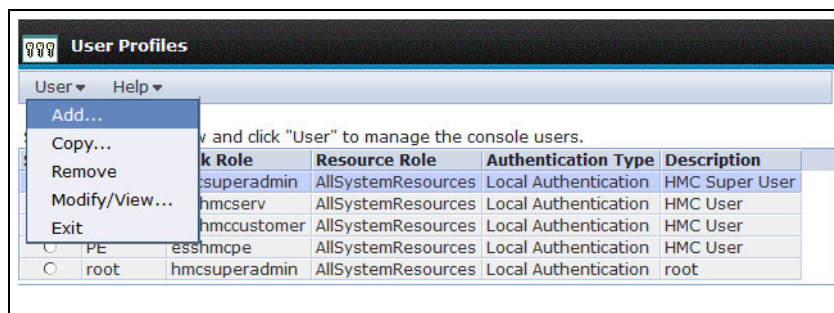


Figure 6-21 User Profiles option menu

2. Click **Add**. The Add User window opens, as shown in Figure 6-22.

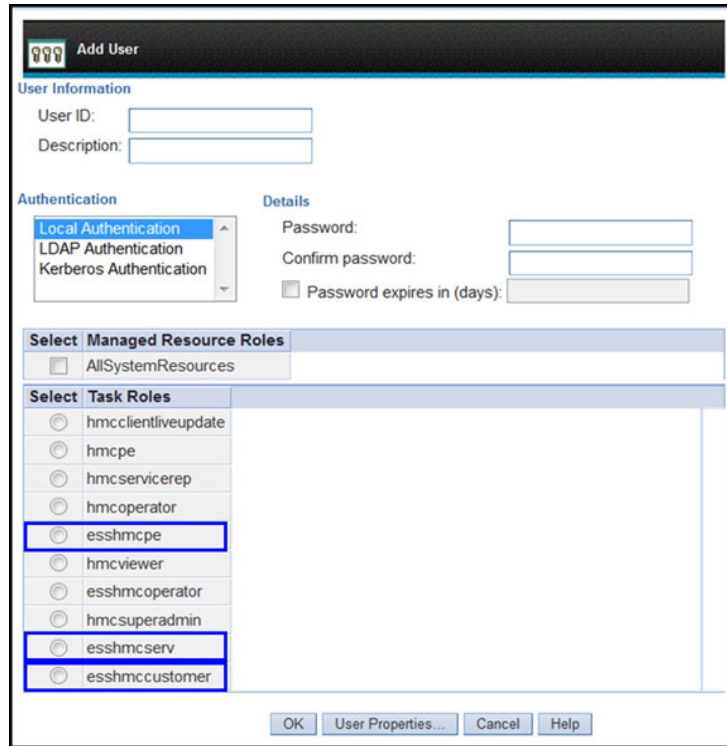


Figure 6-22 Add User window

Only those roles that are outlined by the boxes are valid Task Roles.

3. Complete the following fields:
 - a. Under **Description**, define a user or use HMC User as an example.
 - b. Passwords must adhere to the DS8900F password policies. For more information, see , “After you issue that command, you can run other commands.” on page 187.
 - c. Choose the type of **Authentication** that you want.
 - d. Select **AllSystemResources**, under **Managed Resource Roles**.
 - e. Select the Task Role type.
4. Click **User Properties** to optionally add Timeout and Inactivity values. Ensure that **Allow access via the web** is selected if web access is needed.

Figure 6-23 shows these settings.

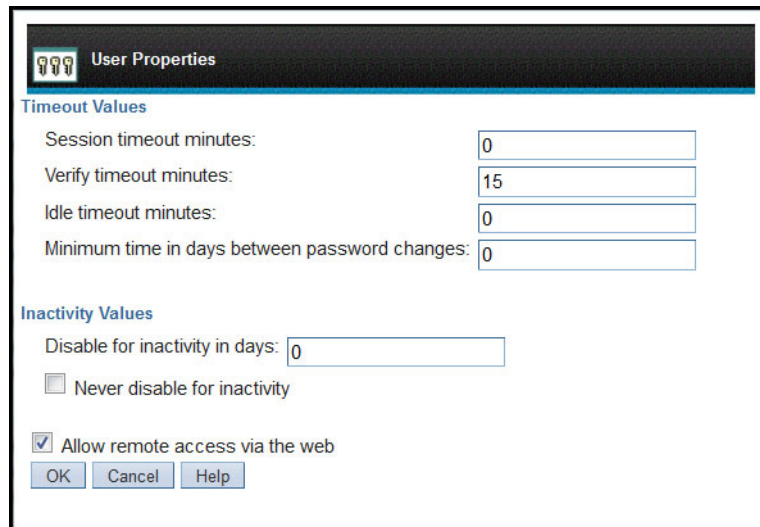


Figure 6-23 User Properties menu

5. Click **OK**, and then click **OK** again. This task is complete.

The User Profiles are updated and list the new user ID. As an example, user ID IBM_RSC was created and is shown in Figure 6-24 and Figure 6-25 on page 195.

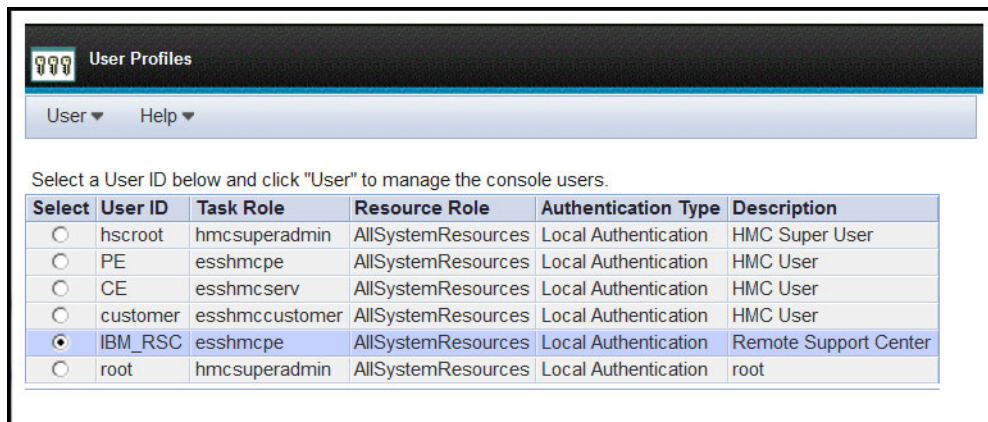


Figure 6-24 IBM_RSC user ID

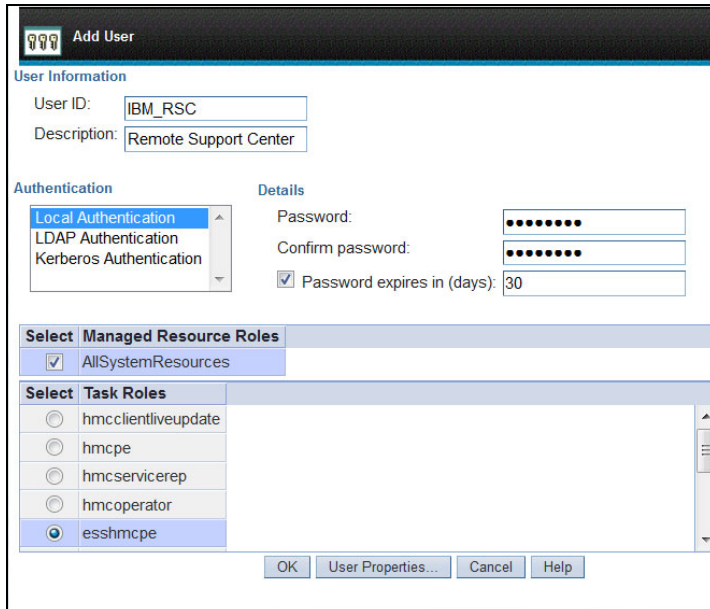


Figure 6-25 IBM_RSC user ID properties

6.5.4 Service Management Console LDAP authentication

Before LDAP authentication is selected, the HMC must first be configured to access an LDAP server. Complete the following steps:

1. In the HMC Management window, select **Configure LDAP**, as shown in Figure 6-26.

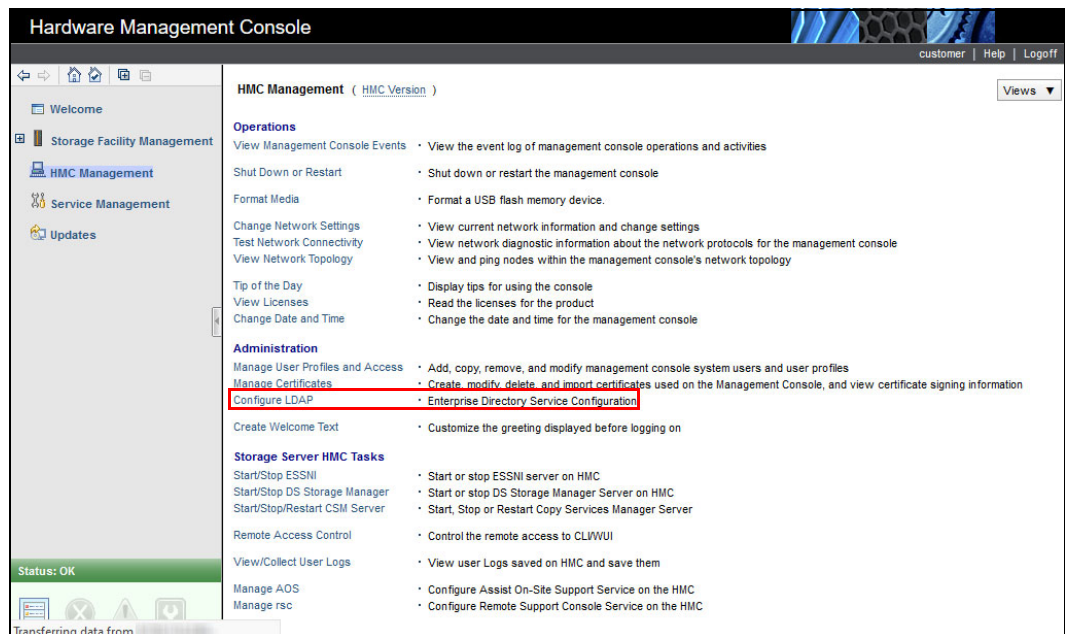


Figure 6-26 Configuring LDAP

2. The window that is shown in Figure 6-27 opens. Perform the following actions:
 - a. Select the **Enable LDAP** box.
 - b. Provide the **Primary URI**, and optionally, the **Backup URI**.
 - c. Choose either **TLS Encryption** or **Non-Anonymous Binding**.
 - d. Select an attribute for **Use the Following Attribute for User Login**.
 - e. Identify the Distinguished Name Tree for search functions.
 - f. Do *not* select **Enable LDAP for Remote User Management**.
3. Click **OK** to complete this task.




Figure 6-27 LDAP Server Definition window

6.5.5 Multifactor authentication (MFA)

Multifactor authentication (MFA) is a stronger method of authentication than just a password. It is typically used to protect sensitive data. MFA requires multiple proofs of identification to gain access.

Starting with DS8000 Release 9.3.2 MFA enables the DS8000 storage system user to configure remote authentication with RSA SecurID Authentication Manager or with direct LDAP+RSA SecurID Authentication Manager. It supports a PIN® (something that the user knows) and a token (something that the user has) as factors of authentication.

You can also configure IBM RACF® to use MFA by implementing IBM Z Multi-Factor Authentication (IBM Z MFA).

For more information, see *LDAP Authentication for IBM Storage DS8000 Systems: Updated for DS8000 Release 9.3.2*, REDP-5460.

6.6 Secondary Management Console

The secondary MC is used for redundancy and it is part of the DS8900F ME. The primary MC is referred to as *MC1*, and the secondary MC is referred to as *MC2*. The two MCs run in a dual-active configuration, so either MC can be used at any time. Each MC is assigned a role of either primary (normally MC1) or secondary (normally MC2). Certain service functions can be performed only on the primary MC.

The DS8900F can run all storage duties while the MC is down or offline, but configuration, error reporting, and maintenance capabilities become severely restricted. Any organization with high availability (HA) requirements should strongly consider deploying an MC redundant configuration.

Important: The primary and secondary MCs are not available to be used as general-purpose computing resources.

6.6.1 Management Console redundancy benefits

MC redundancy provides the following advantages:

- ▶ Enhanced maintenance capability
Because the MC is the only interface that is available for service personnel, an alternative MC provides maintenance operational capabilities if the internal MC fails.
- ▶ Greater availability for power management
Using the MC is the only way to safely power on or power off the DS8900F. The secondary MC is necessary to shut down the DS8900F if the primary MC fails.
- ▶ Greater availability of encryption deadlock recovery
If the DS8900F is configured for Full Disk Encryption (FDE) and an encryption deadlock situation occurs, the use of the MC is the only way to input a recovery key to allow the DS8900F to become operational.
- ▶ Greater availability for Advanced CS
Because all CS functions are driven by the MC, any environment that uses Advanced CS must include dual MCs for operational continuity.
- ▶ Greater availability for configuration operations
All configuration commands must go through the MC. This requirement is true regardless of whether access is through IBM Spectrum Control, DS CLI, or DS GUI. The secondary MC allows these operations to continue if the primary MC fails.

When a configuration or CS command is run, the DS CLI or DS GUI sends the command to the first MC. If the first MC is unavailable, it automatically sends the command to the second MC instead. Typically, you do not need to reissue the command.

Any changes that are made by using one MC are instantly reflected in the other MC. No host data is cached within the MC, so no cache coherency issues occur.



IBM DS8900F features and licensed functions

This chapter describes licensed functions and their activation for the IBM DS8900F.

This chapter covers the following topics:

- ▶ DS8900F licensed functions
- ▶ Activating licensed functions
- ▶ Licensed scope considerations

7.1 DS8900F licensed functions

The licensed functions are bundled into groups, as shown in Table 7-1.

Table 7-1 IBM DS8000 licensed functions

Licensed function for DS8000 with Enterprise Choice warranty	License scope	IBM 9031-FF8 indicator Feature Code numbers
Base Function (BF)	ALL	8151 - 8160
Copy Services (CS)	ALL, Fixed-Block (FB), or Count Key Data (CKD)	8250 - 8260
Z Synergy Services (zsS)	CKD	8350 - 8360
CS on Hardware Management Console (HMC) ^a		8451

a. The license for IBM Copy Services Manager on the HMC server must be purchased as a separate software license.

The IBM Copy Services Manager provides an advanced GUI to easily and efficiently manage CS. IBM Copy Services Manager is available on the DS8000 HMC, which eliminates the need to maintain a separate server for CS functions. For that reason, and in addition to the other license bundles that are shown in Table 7-1, the IBM Copy Services Manager for HMC license can be configured along with these bundles, and it is enabled by using a data storage feature activation (DSFA) key. IBM Copy Services Manager enablement files are activated on the HMC when the key is applied.

The grouping of licensed functions facilitates ordering, which differs from earlier DS8000 models for which licensed functions were more granular and ordered specifically.

The four license bundles contain the following functions:

- ▶ BF license:
 - Operating Environment License (OEL)
 - Logical configuration support for FB (open systems)
 - Thin provisioning
 - Easy Tier
 - Encryption authorization
- ▶ CS license:
 - FlashCopy
 - Metro Mirror (MM)
 - Global Mirror (GM)
 - Metro/Global Mirror (MGM)
 - z/Global Mirror
 - z/Global Mirror Resync
 - Multi-Target Peer-to-Peer Remote Copy (PPRC)
 - Safeguarded Copy

- ▶ zsS:
 - Fibre Channel connection (IBM FICON) attachment
 - Parallel access volume (PAV)
 - HyperPAV
 - SuperPAV
 - High-Performance FICON for IBM Z (zHPF)
 - IBM z/OS Distributed Data Backup (zDDB)
 - Transparent Cloud Tiering (TCT)
 - zHyperLink
- ▶ IBM Copy Services Manager on the HMC license

IBM Copy Services Manager facilitates the use and management of CS functions, such as the remote mirror and copy functions (MM and GM) and the point-in-time copy (PTC) function (FlashCopy). IBM Copy Services Manager is available on the HMC, which eliminates the need to maintain a separate server for CS functions.

Licensed functions enable the operating system and functions of the storage system. Some features, such as the operating system, are always enabled, and other functions are optional. Licensed functions are purchased as 5341 machine function authorizations for billing purposes.

Each licensed function indicator feature that is ordered enables that function at the system level, and it requires an equivalent 9031 function authorization. The licensed function indicators are also used for maintenance billing purposes.

Starting with DS8900F R9.3, maintenance and support fall under Expert Care, which defines the support duration (1, 2, 3, 4, or 5 years) and the service level (Advanced or Premium). When purchasing IBM DS8900F (machine type 5341), the inclusion of Expert Care is mandatory. For more information, see 7.4, “Expert Care” on page 220.

- ▶ All DS8900F models are sold with a 1-year warranty. This warranty is extended by Expert Care from 2 to 5 years. The machine type 5341 no longer indicates the warranty period.
- ▶ Each license function authorization is associated with a fixed 1-year function authorization 9031-FF8.
- ▶ The licensed function indicator feature numbers enable the technical activation of the function, subject to a feature activation code that is made available by IBM, which must be applied by the client.
- ▶ Licensed functions are activated and enforced with a defined license scope. *License scope* refers to the type of storage and the type of servers that the function can be used with. For example, the zsS licenses are available only with the CKD (z/FICON) scope.

The BFs are mandatory. The BFs must always be configured for both mainframe and open systems, which have a scope of ALL. Also, to configure CKD volumes, Feature Code 8300 is required.

With CS, if these services are used only for either mainframe *or* open systems, the restriction to either FB or CKD is possible. However, most clients want to configure CS for the scope ALL.

For each group of licensed functions, specific Feature Code numbers indicate the licensed capacity, as shown in Table 7-2. These Feature Codes vary depending on model. For more information, see “Ordering granularity” on page 206.

Table 7-2 License Feature Codes example

Feature Code			Feature Code for licensed function indicator for raw capacity
BF	CS	zsS	
8151	8251	8351	10 TB (up to 100 TB capacity)
8152	8252	8352	15 TB (100.1 TB - 250 TB capacity)
8153	8253	8353	25 TB (250.1 TB - 500 TB capacity)
8154	8254	8354	75 TB (500.1 - 1,250 TB capacity)
8155	8255	8355	175 TB (1,250.1 TB - 3,000 TB capacity)
8156	8256	8356	300 TB (3,000.1 TB - 6,000 TB capacity)
8160	8260	8360	500 TB (6,000.1 TB - 12,000 TB capacity)

7.1.1 General introduction to licensing

Several of the orderable Feature Codes must be activated through the installation of a corresponding license key bundle. These feature bundle codes are listed in Table 7-1 on page 200. Certain features can be configured directly for the client by the IBM marketing representative during the ordering process.

Ordering features that require a license key

Important: All CSs are bundled.

The following features are available after the license bundle is activated:

- ▶ MM is a synchronous way to perform remote replication. GM enables asynchronous replication, which is useful for longer distances and lower bandwidth.
- ▶ MGM enables cascaded 3-site replication, which combines synchronous mirroring to an intermediate site with asynchronous mirroring from that intermediate site to a third site at a long distance.

Combinations with other CS features are possible and sometimes needed. Usually, the 3-site MGM installation also requires an MM sublicense on site A with the MGM license (and even a GM sublicense, if after a site B breakdown you want to resynchronize site A and site C). At site B, on top of the MGM, you also need the MM and GM licenses. At site C, you then need sublicenses for MGM, GM, and FlashCopy.

- ▶ Multiple-Target PPRC (MT-PPRC) enhances disaster recovery (DR) solutions by allowing data at a single primary site to be mirrored to two remote sites simultaneously. The function builds and extends MM and GM capabilities and is supported on DS8900F, DS8880, on later DS8870 firmware, and on IBM Copy Services Manager or IBM Z software, such as IBM Geographically Dispersed Parallel Sysplex (IBM GDPS) / MTMM.

Various interfaces and operating systems (OSs) support the function. For the DS8900F family, this feature is integrated with the CS license bundle.

- ▶ Two possibilities exist for FlashCopy PTC: Use it with thick (standard) volumes or thin-provisioned extent space efficient (ESE) volumes.
The ESE thin volumes can also be used in remote mirroring relationships. ESE volumes offer the same good performance as standard (thick) volumes, and can be managed by IBM Easy Tier.
- ▶ Safeguarded Copy enables you to create snapshots for Logical Corruption Protection (LCP). It provides many recovery points from which to restore data in case of logical corruption or destruction of data. If Safeguarded Copy is used, you should additionally mark the Feature Code 0785 indicator option when ordering.
- ▶ The z/OS Global Mirror (zGM) license, which is also known as Extended Remote Copy (XRC), enables z/OS clients to copy data by using System Data Mover (SDM). This copy is asynchronous.
- ▶ As previously noted, IBM Copy Services Manager on HMC offers full IBM Copy Services Manager functions and must be enabled by using a DSFA activation key. IBM Copy Services Manager enablement files are activated on the HMC when the key is applied. The license for IBM Copy Services Manager on the HMC server must be purchased as a separate software license.
- ▶ For IBM Z clients, PAVs allow multiple concurrent I/O streams to the same CKD volume. HyperPAV reassigns the alias addresses dynamically to the base addresses of the volumes based on the needs of a dynamically changing workload. Both features result in such large performance gains that for many years they were configured as an effective standard for mainframe clients, similar to FICON, which is required for z/OS.
SuperPAV is an extension to HyperPAV support and allows aliases to be borrowed from eligible peer logical control units (LCUs).
- ▶ zHPF is a feature that uses a protocol extension for FICON and allows data for multiple commands to be grouped in a single data transfer. This grouping increases the channel throughput for many workload profiles. It works on all newer IBM zEnterprise Systems and it is preferred for these systems because of the performance gains that it offers.
- ▶ zDDB is a feature for clients with a mixture of mainframe and distributed workloads to use their powerful IBM Z host facilities to back up and restore open systems data. For more information, see *IBM System Storage DS8000: z/OS Distributed Data Backup*, REDP-4701.
- ▶ Easy Tier is available in the following modes:
 - Automatic mode works on the subvolume level (extent level) and allows auto-tiering in hybrid extent pools. The most-accessed volume parts go to the upper tiers. In single-tier pools, it allows auto-rebalancing if it is turned on.
 - Manual Dynamic Volume Relocation (DVR) mode works on the level of full volumes and allows volumes to be relocated or restriped to other places in the DS8000 online. It also allows ranks to be moved out of pools. For more information, see *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667.
 - Easy Tier Heat Map Transfer (HMT) automatically replicates a heat map to remote systems to ensure that they are also optimized for performance and cost after a planned or unplanned outage. For more information, see *IBM DS8870 Easy Tier Heat Map Transfer*, REDP-5015.
- ▶ The Encryption Authorization feature provides data encryption by using IBM Full Disk Encryption (FDE) and key managers, such as IBM Security Guardium Key Lifecycle Manager. The key manager must be licensed separately.

For more information about these features, see *IBM DS8900F Introduction and Planning Guide*, SC27-9560.

7.1.2 Licensing cost structure

For the three possible license bundles, the following general guidelines apply:

- ▶ BF license

The license must be at least equal to the total raw capacity of the storage system, which includes the raw capacity of any expansion frames. Select the full raw capacity and order the number of TBs to support the total raw capacity of your storage system.

- ▶ CS license

The license can be equal to, less than, or more than the total raw capacity of the storage system. You must license the terabytes that you use (provisioned capacity) in copy relationships.

- ▶ zsS license

The license can be equal to or less than the total raw capacity of the storage system, but order less only if you have a mixed machine (FB and open, and CKD and mainframe). You must license the full CKD / IBM Z raw capacity for zsS, that is, the raw capacity of all ranks that will be formatted as CKD ranks.

Copy Services license specifics

You can order the CS license to support the *total provisioned capacity* of all volumes that are involved in one or more CS functions. However, this sort of subcapacity licensing (less than the total raw capacity) requires capacity monitoring and a steady remote connection on the client side.

Important: With the CS license bundle, order subcapacity licensing, which is less than the total physical raw capacity, only when a steady remote connection for the DS8000 is available.

By using a remote connection for Call Home, the CS license can be based on the usable capacity of the volumes that will potentially be in CS relationships. This amount typically is less than the total raw capacity.

Note: The CS license goes by the capacity of all volumes that are involved in at least one CS relationship. The CS license is based on the provisioned capacity of volumes and not on raw capacity. If overprovisioning is used on the DS8000 with a significant number of CS functions, the CS license needs to be equal only to the total provisioned capacity. This situation is true even if the logical volume capacity of volumes in CS is greater.

For example, with overprovisioning, if the total rank raw capacity of a DS8900F is 100 TB but 200 TB of thin-provisioning volumes are in MM, only a 100 TB of CS license is needed.

For FlashCopy volumes, you must count the source plus target volumes as provisioned capacity. Several examples are shown in “Pricing examples for Copy Services” on page 205.

Pricing examples for Copy Services

The following examples are provided to illustrate your CS licensing requirements in an FB environment:

- ▶ Scenario 1: For FlashCopy for a 10 TB source, the purchase of a 20 TB capacity CS license is required.
- ▶ Scenario 2: To use MM on a 10 TB source and then FlashCopy on a 10 TB target, the purchase of a 10 TBs CS license on the source and a 20 TB CS license on the target DS8000 is required.
- ▶ Scenario 3: To use GM on a 10 TB source and then FlashCopy on a 10 TB target DS8000, the purchase of a 10 TB CS license on the source and a 20 TB CS license on the target DS8000 is required.
- ▶ Scenario 4: To use MGM on a 10 TB source and then FlashCopy on a 10 TB target, the purchase of a 10 TB CS license on the source and secondary, and the purchase of a 20 TB CS license on the target is required.

However, consider that with MGM, certain scenarios can require more FlashCopy targets on the local machines, and so larger CS terabyte scopes are necessary.

- ▶ Scenario 5: A client wants to perform GM for a 10 TB source and use FlashCopy on the target for practicing DR, but they do not want to affect the normal GM. This situation requires a GM secondary, GM Journal, and a FlashCopy volume on the secondary system. The source DS8900F requires a 10 TB CS license, and the target DS8880 requires a 30 TB CS license.
- ▶ Scenario 6: To perform 4-site replications, the purchase of the correct capacity license requirement for each storage system is required.

Z Synergy Services licensing

A zsS license is required for only the total physical capacity that is logically configured as CKD ranks for use with IBM Z host systems.

Note: If zDDB is used on a system with no CKD ranks, a 10 TB zsS license must be ordered to enable the FICON attachment functions.

Drive features

The BF is based on the raw (decimal terabyte) capacity of the drives. The pricing is based on the drive performance, capacity, and other characteristics that provide more flexible and optimal price and performance configurations.

To calculate the raw (gross) physical capacity, multiply for each drive set the number of drives with their individual capacities. Therefore, for example, a drive set of sixteen 3.84 TB drives has a 61.44 TB raw capacity.

Table 7-3 shows the Feature Codes for high-performance flash drive sets.

Table 7-3 Feature Codes for high-performance flash drive sets

Feature Code	Drive capacity	Drive type	Drive speed	Encryption-capable	Drives per set
1611	800 GB	2.5-inch flash	N/A	Yes	16
1612	1.6 TB	2.5-inch flash	N/A	Yes	16
1613	3.2 TB	2.5-inch flash	N/A	Yes	16

Table 7-4 shows the Feature Codes for high-capacity flash drive sets.

Table 7-4 Feature Codes for high-capacity flash drive sets

Feature Code	Drive capacity	Drive type	Drive speed	Encryption-capable	Drives per set
1622	1.92 TB	2.5-inch flash	N/A	Yes	16
1623	3.84 TB	2.5-inch flash	N/A	Yes	16
1624	7.68 TB	2.5-inch flash	N/A	Yes	16
1625	15.36 TB	2.5-inch flash	N/A	Yes	16

Important: Check with an IBM Systems Service Representative (IBM SSR) or go to the IBM website for an up-to-date list of available drive types.

Related information: New storage system expansions for DS8900F are delivered only with FDE drives.

Easy Tier is a license feature that is available at no charge. Therefore, it is configured, by default, with the BF license bundle.

The Database Protection (for open and FB) feature and the Thin Provisioning feature come with the BF license bundle.

The zDDB feature comes with the zsS license bundle.

Ordering granularity

You order the license bundles by the terabyte, but not by a single terabyte. The granularity is slightly larger. For example, below 100 TB total raw capacity, the granularity increment for an upgrade is 10 TB. With larger total capacities, the granularity is larger. For more information, see Table 7-5.

Table 7-5 Ordering granularity

Tier	Minimum (TB)	Maximum (TB)	Tier TB granularity	Feature quantity (maximum)	Range delta (TB)	Feature number
1	1	100	10	10	100	8x51
2	101	250	15	10	150	8x52
3	251	500	25	10	250	8x53
4	501	1250	75	10	750	8x54
5	1251	3000	175	10	1750	8x55
6	3001	6000	300	10	3000	8x56
7	6001	12000	500	12	6000	8x60

Tip: For more information about the features and considerations when you order DS8900F licensed functions, go to [IBM Offering Information](#) and search for the *IBM DS8900F Models 993, 994, 996, 998, and E96* announcement letter by using the *DS8900F* keyword as a search term.

7.2 Activating licensed functions

You can activate the license keys of the DS8000 after the IBM SSR completes the storage complex installation. If you plan to use the Storage Management GUI to configure your new storage, after the initial login as `admin`, the setup wizard guides you to download your keys from the DSFA website and activate them. However, if you plan to use the Data Storage Command-line Interface (DS CLI) to configure your new storage, you must first obtain the necessary keys from the [DSFA website](#).

Before you connect to the DSFA website to obtain your feature activation codes, ensure that you have the following items:

- ▶ The IBM License Function Authorization documents. If you are activating codes for a new storage unit, these documents are included in the shipment of the storage unit. If you are activating codes for an existing storage unit, IBM sends the documents to you in an envelope.
- ▶ A USB memory device that can be used for downloading your activation codes if you cannot access the DS Storage Manager from the system that you are using to access the DSFA website. Instead of downloading the activation codes in softcopy format, you can print the activation codes and manually enter them by using the DS Storage Manager GUI or the DS CLI. However, this process is slow and error-prone because the activation keys are 32-character strings.

7.2.1 Obtaining DS8000 machine information and activating license keys

To obtain the license activation keys from the DSFA website, you must know the serial number and machine signature of your DS8000 unit.

You can obtain the required information by using the DS Storage Management GUI or the DS CLI. If you use the Storage Management GUI, you can obtain and apply your activation keys at the same time. These options are described next.

DS Storage Management GUI

To obtain the required information by using the DS Storage Management GUI, complete these steps:

1. Open a browser and enter `https://< IP address of HMC >`.
2. Log in by using a user ID with administrator access. If you are accessing the system for the first time, contact your IBM SSR for the user ID and password. After a successful login, the system monitor window opens.

- If this machine is new, a System Setup wizard window opens automatically that guides you through the initial setup and configuration tasks, as shown in Figure 7-1.

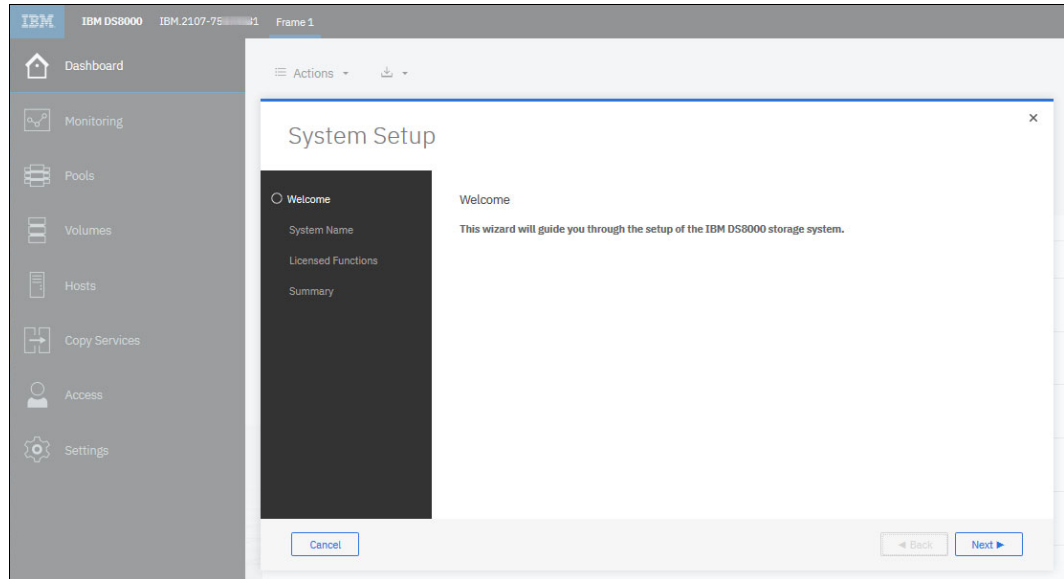


Figure 7-1 System Setup wizard including Licensed Functions activation

Note: Before you begin this task, resolve any current DS8000 problems that might exist. You can contact IBM Support to help you resolve these problems.

- To begin the guided procedure to acquire and activate your feature activation keys, select **System Setup** → **Licensed Functions**, and then complete the **Activate Licensed Functions** routine, as shown in Figure 7-2.

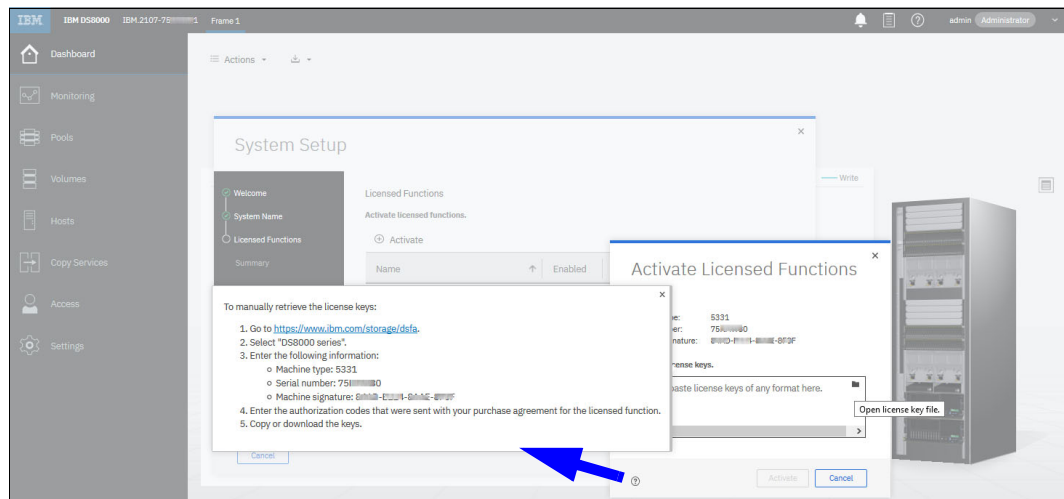


Figure 7-2 Entering the license keys with help

Note: You can download the keys and save the XML file to the folder that is shown here, or you can copy the license keys from the IBM DSFA website.

- After you enter all your license keys, click **Activate** to start the activation process, as shown in Figure 7-3.

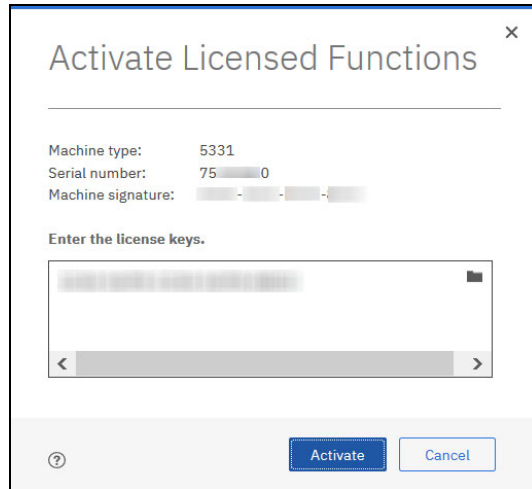


Figure 7-3 Adding the license keys

- Click **Summary** in the System Setup wizard to view the list of licensed functions or feature keys that are installed on your DS8000, as shown in Figure 7-4.

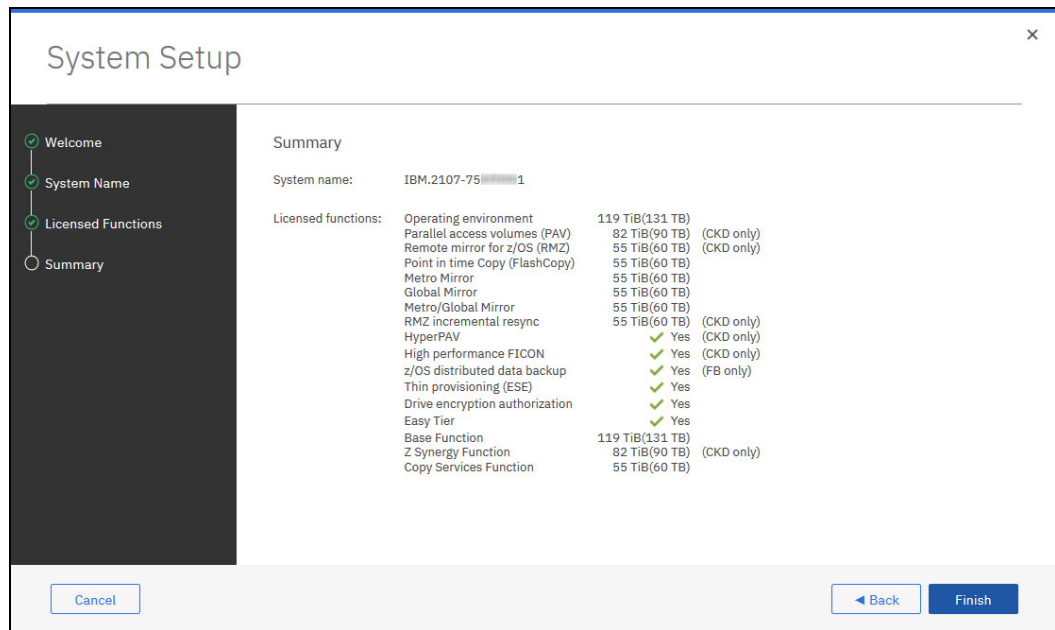


Figure 7-4 Summary of licensed functions

- If you must activate more feature keys after the initial installation, click the **Settings** icon in the left menu, select **Licensed Functions**, and click the **Activate** icon, which opens another window where you enter the keys, as shown in Figure 7-5.

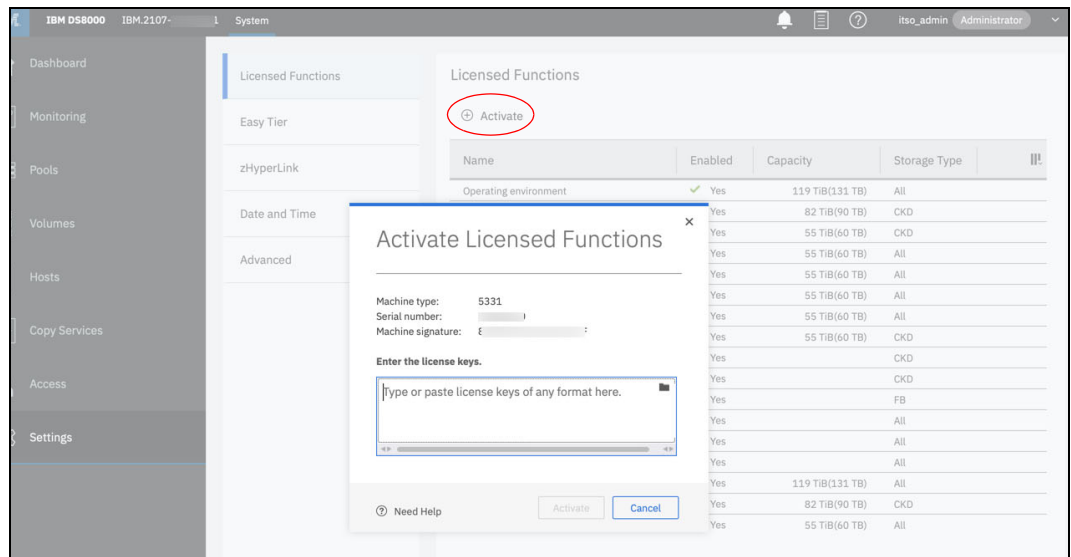


Figure 7-5 Licensed Functions

- To obtain the machine signature and machine type and model (MTM) after the installation, go to the Dashboard and click **Actions** → **Properties**, as shown in Figure 7-6.

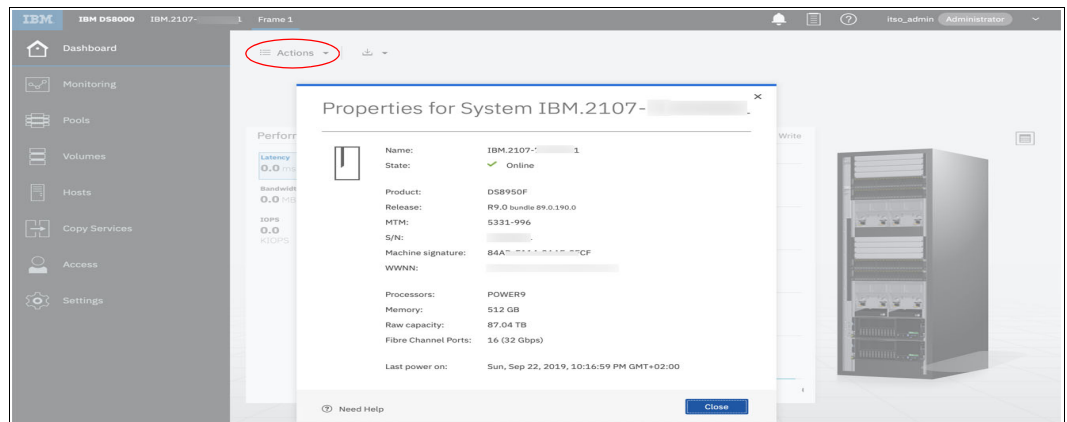


Figure 7-6 Properties window showing the machine signature and MTM

Important: The initial enablement of any optional DS8000 licensed function is a concurrent activity (assuming that the correct level of Licensed Internal Code (LIC) is installed on the system for the function).

The following activation activities are disruptive and require an initial machine load or restart of the affected image:

- ▶ Removal of a DS8000 licensed function to deactivate the function.
- ▶ A lateral change or reduction in the license scope. A *lateral change* is defined as changing the license scope from FB to 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.

Note: Before you begin this task, you must resolve any current DS8000 problems that exist. You can contact IBM Support for help with resolving these problems.

9. Click **Activate** to enter and activate your licensed keys, as shown in Figure 7-2 on page 208.
10. Wait for the activation process to complete and select **Licensed Functions** to show the list of activated features.

DS Command-Line Interface

To obtain the required information by using the DS CLI, log on to the DS CLI and run the **lssi** and **shows** commands, as shown in Figure 7-7.

```
dsccli> lssi
Date/Time: 01 April 2022 15:16:29 CEST IBM DSCLI Version: 7.9.21.80 DS: -
Name      ID              Storage Unit    Model WNNN          State  ESSNet
=====
ds8k-r9-01 IBM.2107-75HAL91 IBM.2107-75HAL90 994   5005076312345678 Online Enabled

dsccli> showsi
Date/Time: 01 April 2022 15:16:35 CEST IBM DSCLI Version: 7.9.21.80 DS: -
Name          ds8k-r9-01
desc          Sand Shark
ID            IBM.2107-75HAL91
Storage Unit  IBM.2107-75HAL90
Model        994
WWNN         5005076312345678
Signature     abcd-ef12-3456-7890
State        Online
ESSNet       Enabled
Volume Group V0
os400Serial   050
NVS Memory   8.0 GB
Cache Memory 143.1 GB
Processor Memory 183.9 GB
MTS          IBM.5331-75HAL90
numegsupported 16
ETAutoMode   all
ETMonitor    all
IOPMmode     Disabled
ETCCMode     -
ETHMTMode    Enabled
ETSRMode     Enabled
ETTierOrder  High performance
ETAutoModeAccel Disabled
```

Figure 7-7 Obtaining DS8000 information by using the DS CLI

Note: The **shows** command can take the storage facility image (SFI) serial number as a possible argument. The SFI serial number is identical to the storage unit serial number, except that the SFI serial number ends in 1 instead of 0 (zero).

Gather the following information about your storage unit:

- ▶ The Machine Type - Serial Number (MTS), which is a string that contains the machine type and the serial number. The machine type, now mostly 5341, here above is 5331, and the last 7 characters of the string are the machine's serial number (XYABCDE), which always ends with 0 (zero).
- ▶ The model, which, for example, is 996 for a DS8950F.
- ▶ The machine signature, which is found in the Machine signature field and uses the following format:

ABCD-EF12-3456-7890

Table 7-6 documents this information, which is entered at the DSFA website to retrieve the activation codes.

Table 7-6 DS8000 machine information

Property	Your storage unit's information
Machine type	
Machine's serial number	
Machine signature	

7.2.2 Obtaining the activation codes

If you plan to use the DS CLI to configure your system, you must obtain your activation keys before you configure your machine.

Note: A DS8880 is shown in the following figures. However, the steps are identical for all models of the DS8000 family.

To obtain the activation codes, complete the following steps:

1. As shown in Figure 7-8, connect to the [DSFA website](#).

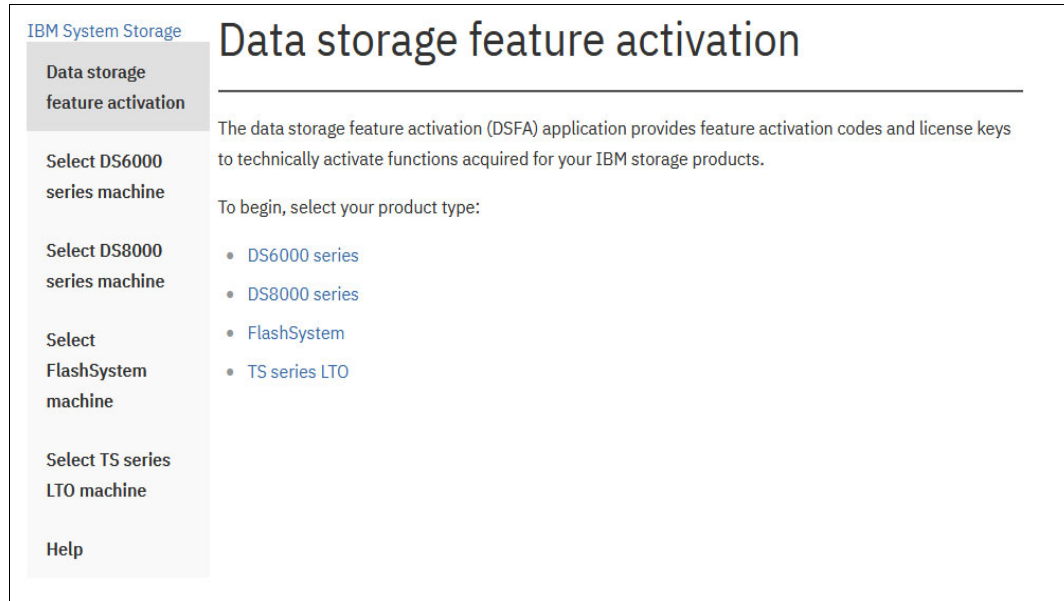


Figure 7-8 DSFA website

2. Click **DS8000 series**. The Select DS8000 series machine window opens, as shown in Figure 7-9. Select the appropriate 5341, 533x, 283x, 242x, or 958x machine type.

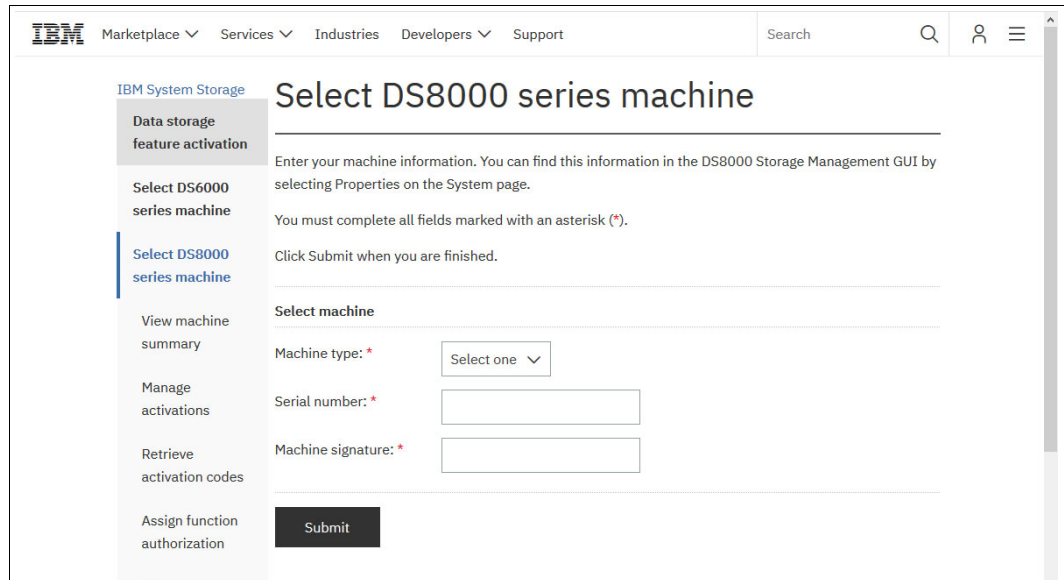


Figure 7-9 DS8000 DSFA machine information entry window

3. Enter the machine information that was collected in Table 7-6 on page 213 and click **Submit**. The View machine summary window opens, as shown in Figure 7-10.

The screenshot shows the IBM DSFA 'View machine summary' window. The page header includes the IBM logo and navigation links: 'Industries & solutions', 'Services', 'Products', 'Support & downloads', and 'My IBM'. A search bar is also present. The breadcrumb trail indicates the user is in 'IBM System Storage'. The sidebar on the left contains several menu items, with 'View machine summary' highlighted. The main content area displays the machine details: 'IBM 2831 Model 980' and 'Serial number 75-A44400'. Below this, there are instructions on how to obtain feature activation codes, a list of three steps, and a note that the page displays information for functions present on the machine. A 'Miscellaneous' section contains a table with one entry: Feature '8300' and Description 'FICON indicator'. At the bottom, there is a license summary table for 'IBM 2836 Model LF8 Serial number 75+66666'. The license table has four columns: 'Description', 'Total license', 'Assigned', and 'Unassigned'. The rows are: 'Base function' (325.0 TB total, 0.0 TB assigned, 325.0 TB unassigned), 'z-synergy services' (115.0 TB total, 0.0 TB assigned, 115.0 TB unassigned), and 'Copy services' (275.0 TB total, 0.0 TB assigned, 275.0 TB unassigned).

Figure 7-10 DSFA View machine summary window

The View machine summary window shows the total purchased licenses and the number of licenses that are currently assigned. When you assign licenses for the first time, the Assigned field shows 0.0 TB.

4. On the left, click **Manage activations**. The Manage activations window opens, as shown in Figure 7-11. For each license type and storage image, enter the following information that is assigned to the storage image:
 - Select the license scope from the list box:
 - FB (Fixed-Block data)
 - CKD
 - All
 - Type the capacity value (in TB) to assign to the storage image.

The capacity values are expressed in decimal terabytes. The sum of the storage image capacity values for a license cannot exceed the total license value.

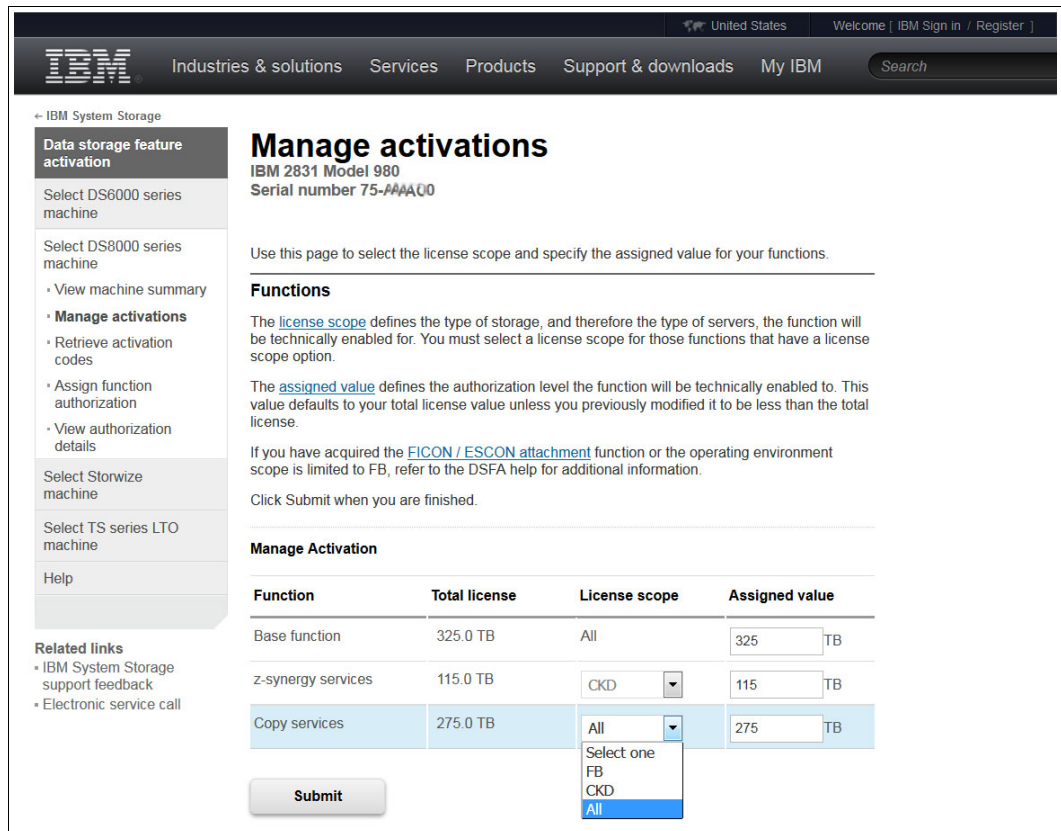


Figure 7-11 DSFA Manage activations window

5. After the values are entered, click **Submit**.
6. On the left on Figure 7-11, select **Retrieve activation codes**.

- The Retrieve activation codes window opens, which shows the license activation codes for the storage image, as shown in Figure 7-12. Print the activation codes or click **Download now** to save the activation codes in an XML file that you can import into the DS8000.

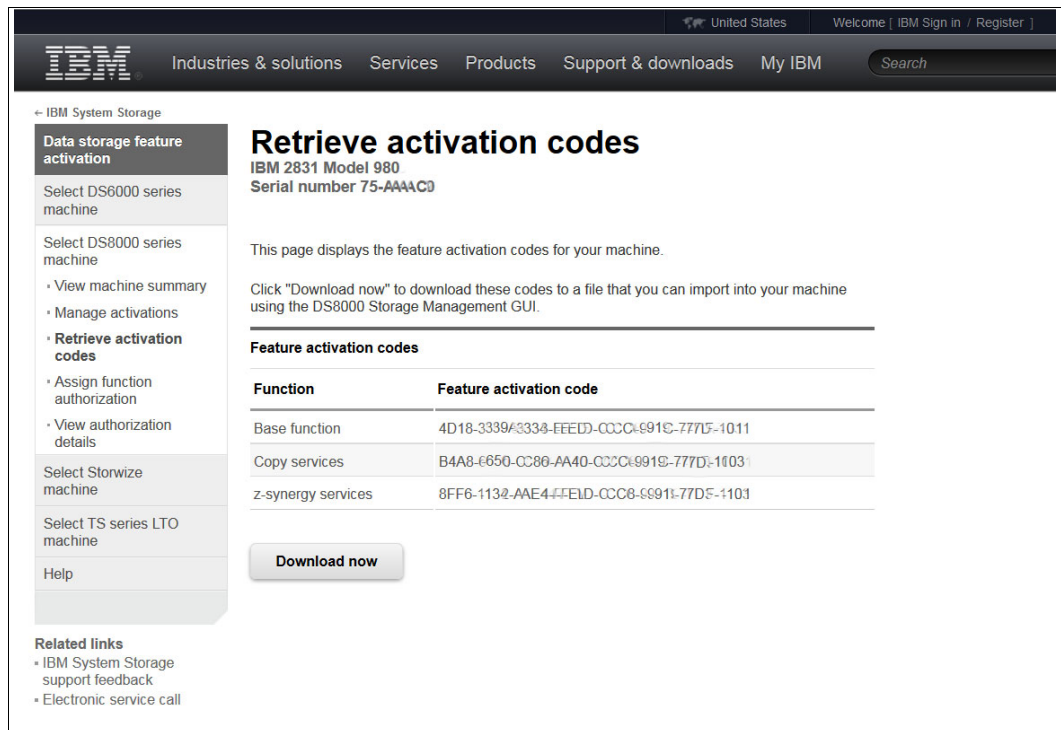


Figure 7-12 DSFA Retrieve activation codes window

- Click **Settings** → **System** → **Licensed Functions** → **Activate** to enter the activation codes after the initial installation. Click the black symbol to import the XML file, or alternatively, you can enter individual license keys, as shown in Figure 7-13.

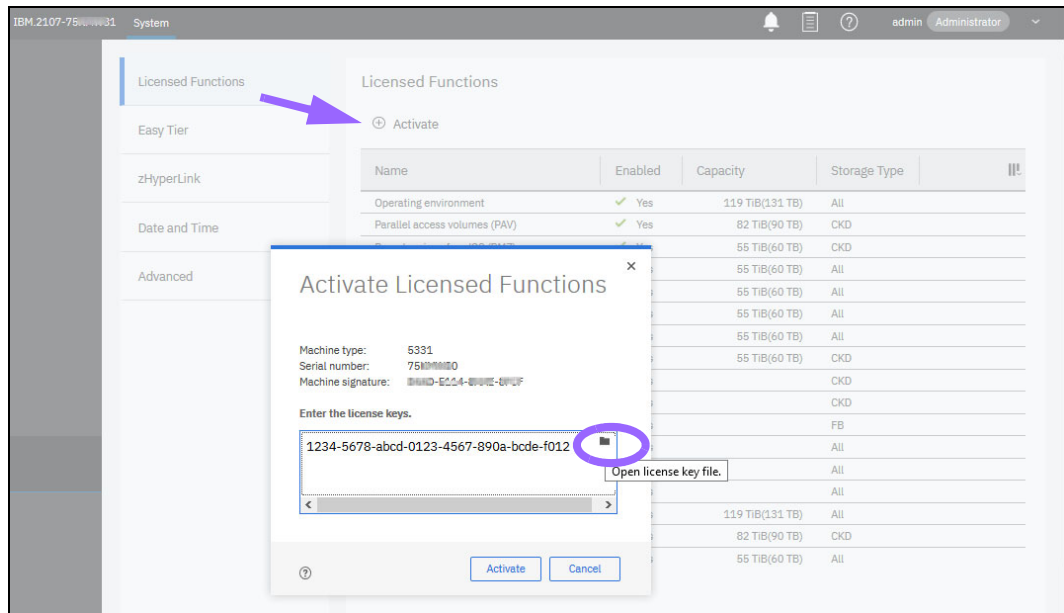


Figure 7-13 Activating more licenses in the GUI

Important: In most situations, the DSFA application can locate your 9031 (or 904y, y=6...9) licensed function authorization record when you enter the DS8900F 5341 (or 533x, x=1...4) serial number and signature. However, if the 90yy licensed function authorization record is not attached to the 53xx record, you must assign it to the 53xx record by using the **Assign function authorization** link on the DSFA application. In this case, you need the 90yy serial number (which you can find on the License Function Authorization document).

7.2.3 Applying activation codes by using the DS CLI

The license keys also can be activated by using the DS CLI. This option is available only if the machine OEL was activated and you installed a compatible DS CLI program on your console.

To apply activation codes by using the DS CLI, complete the following steps:

1. Run the **shows i** command to display the DS8000 machine signature, as shown in Figure 7-14.

```
dscli> shows i
Date/Time: 06 April 2022 15:01:47 CEST IBM DSCLI Version: 7.9.30.154 DS: -
Name          ds8k-r9-01
desc          Sand Shark
ID            IBM.2107-75HAL91
Storage Unit  IBM.2107-75HAL90
Model         998
WWNN          5005076312345678
Signature     abcd-ef12-3456-7890
State         Online
ESSNet        Enabled
Volume Group  V0
os400Serial   6DF
NVS Memory    127.5 GB
Cache Memory  4168.4 GB
Processor Memory 4343.4 GB
MTS           IBM.5341-75HAL90
numegsupported 16
ETAutoMode    tiered
ETMonitor     automode
IOPMmode      Disabled
ETCCMode      -
ETHMTMode     Enabled
ETSRMode      Enabled
ETTierOrder   High performance
ETAutoModeAcce1 Disabled
```

Figure 7-14 DS CLI shows i command

2. Obtain your license activation codes from the IBM DSFA website, as described in 7.2.2, “Obtaining the activation codes” on page 213.

- Enter the **applykey** command at the following DS CLI. The **-file** parameter specifies the key file. The second parameter specifies the storage image.

```
dscli> applykey -file c:\53xx_75XXXX0.xml IBM.2107-75XXXX1
```

Or you can apply individual keys by running the following command:

```
dscli> applykey -key f190-1234-1234-1234-5678-1234-5678 IBM.2107-75XXXX1
CMUC00199I applykey: License Machine Code successfully applied to storage image
IBM.2107-75XXXX1.
```

- Verify that the keys were activated for your storage unit by running the **lskey** command, as shown in Figure 7-15.

dscli> lskey		
Activation Key	Authorization Level (TB)	Scope
=====		
Base function	130.8	A11
Copy services	130.8	A11
Encryption Authorization	on	A11
Global Mirror (GM)	130.8	A11
High Performance FICON for System z (zHPF)	on	CKD
IBM HyperPAV	on	CKD
IBM System Storage DS8000 Thin Provisioning	on	A11
IBM System Storage Easy Tier	on	A11
IBM z/OS Distributed Data Backup	on	FB
Metro/Global Mirror (MGM)	130.8	A11
Metro Mirror (MM)	130.8	A11
Operating environment (OEL)	130.8	A11
Parallel access volumes (PAV)	60.5	CKD
Point-in-time copy (PTC)	130.8	A11
RMZ Resync	130.8	CKD
Remote Mirror for z/OS (RMZ)	130.8	CKD
z-synergy services	60.5	CKD

Figure 7-15 Using the **lskey** command to list the installed licenses

For more information about the DS CLI, see *IBM DS8000 Series Command-Line Interface User's Guide*, SC27-9562.

The BF license bundle must be installed before ranks can be formatted for FB (open systems). The zsS license bundle must be installed before ranks can be formatted for CKD (mainframe).

7.3 Licensed scope considerations

An increase in license capacity is concurrent. A deactivation (or decrease in license capacity) takes place only after one more machine initial machine load. Similar considerations apply to the license scopes:

- ▶ The BF license and CS license are available with several license scopes, such as ALL, FB, and CKD.
- ▶ An increase in license scope, for example, changing FB or CKD to ALL, is a concurrent activity.
- ▶ A lateral change, such as changing FB to CKD or changing CKD to FB, or a reduction of the license scope, such as changing ALL to FB or CKD, is a disruptive activity because it requires an initial machine load to activate the change.

Tip: Because the BF license must be ordered for the full physical capacity anyway, and because the CS license can be ordered for only those volumes that are in CS relationships, consider the following tip: For BF and CS, configure these bundles with scope “ALL” from the beginning.

7.4 Expert Care

With Expert Care for IBM DS8900F, customers can easily select their level of service. At the time of hardware purchase, select either the Advanced or Premium tier along with the coverage term, which can be 1 - 5 years.

Expert Care services

IBM Storage Expert Care includes various services, depending on the selected service level. All services are listed in Table 7-7.

Table 7-7 Expert Care services

Services	Advanced	Premium
Hardware Maintenance (IOR = IBM On-site Repair)	24x7 same day	24x7 same day
Support Line (24x7 remote technical support)	Yes	Yes
Predictive Support	Yes	Yes
Storage Insights	Yes	Yes, Pro edition
Technical Account Manager	No	Yes
Enhanced Support Time	No	Yes
Remote Code Load (RCL) (up to 2x per year)	Optional	Yes
On-Site Code Load	Optional	Optional
Media Retention	Optional	Optional

The Technical Account Manager is new role that combines the previous roles of Technical Sales Manager and Technical Advisor. The TAM acts as the single point of contact for the client. They set up a welcome call, schedule monthly activity reports, advise on code currency, help schedule code upgrades, facilitate the initial installation, help with Call Home and remote support setup, and perform other activities.

Enhanced Support Time targets an incident response time of 30 minutes or less for Severity 1 and 2 incidents in the United States and selected countries.

With Predictive Support, IBM pro-actively notifies customers of possible problems to prevent issues from escalating or causing an impact. Predictive Support leverages statistics and performance metrics from IBM Storage Insights. For more information about IBM Storage Insights, see 12.11, “Using IBM Storage Insights” on page 448.

Expert Care Feature Codes

The MTMs for Expert Care are 5131-A0x for Advanced and 5131-P0x for Premium, where x indicates the number of years of support. The indicator for the Expert Care Premium software program number is 577x-ECP.

Table 7-8 shows the Feature Codes for each of the available options.

Table 7-8 Expert Care Feature Codes

Description	Expert Care Advanced 5131 (A01 - A05)	Expert Care Premium 5131 (P01 - P05)
Expert Care Indicator	ALH0	ALH0
1 year	ALK1	ALL1
2 years	ALK2	ALL2
3 years	ALK3	ALL3
4 years	ALK4	ALL4
5 years	ALK5	ALL5
On-Site Code Load	AHY3	AHY2
Remote Code Load	AHY4	Included

The feature codes for Expert Care might differ from region to region. For a full listing, see the relevant IBM Hardware Announcement for your region.

Specific options are also available regarding contact and resolution times, including 1-hour committed contact, 4-hour committed onsite, or 4-, 6-, 8-, 12-, 24-, 48-, or 72-hour committed fix time, each with a corresponding feature code. For more information, contact your IBM Sales Representative.



Part 3

Storage configuration

This part of the book describes the storage configuration tasks.

This part contains the following chapters:

- ▶ Chapter 8, “Configuration flow” on page 225
- ▶ Chapter 9, “IBM DS8900F Storage Management GUI” on page 233
- ▶ Chapter 10, “IBM DS8900F Storage Management Command-line Interface” on page 339



Configuration flow

This chapter provides a brief overview of the sequence of tasks that are required to configure an IBM DS8900F.

This chapter covers the following topics:

- ▶ Configuration worksheets
- ▶ User and role management
- ▶ Encryption
- ▶ Network security
- ▶ Configuration flow
- ▶ General storage configuration guidelines

8.1 Configuration worksheets

Before a new DS8900F system is delivered, it is highly advised to complete the DS8900F configuration worksheets (also called customization worksheets). For more information about these worksheets, including links to download their current version, see Appendix D, “Customization worksheets”, of the *IBM DS8900F Introduction and Planning Guide*, SC27-9560, or see [DS8900 Configuration Worksheets](#). The guide provides detailed information to help you plan a successful installation of the system.

Note: Planning information for all DS8900F models, including the rack-mounted model 993, is covered in the same guide.

The purpose of the configuration worksheets is to enable a smooth installation of the DS8900F system by ensuring that the necessary information is available to the IBM Support Representative during system installation. It is a best practice to present the completed worksheets to the IBM Support Representative before the delivery of the DS8900F system.

The completed customization worksheets specify the initial setup for the following items:

- ▶ **Company information:** Provide important company and contact information. This information is required to ensure that IBM support personnel can reach the appropriate contact person or persons in your organization, or send a technician to service your system in the event of a critical event as quickly as possible.
- ▶ **Hardware Management Console (HMC) network:** Provide the IP address and local area network (LAN) settings. This information is required to establish connectivity to the Management Console (MC).
- ▶ **Remote support, including Call Home:** Provide information to configure Remote Support and Call Home. This information helps to ensure timely support for critical serviceable events on the system.
- ▶ **Notification:** Provide information to receive Simple Network Management Protocol (SNMP) traps and email notifications. This information is required if you want to be notified about serviceable events.
- ▶ **Power control:** Provide your preferences for the power mode on the system.
- ▶ **Control switch:** Provide information to set up the control switches on the system. This information is helpful if you want to customize settings that affect host connectivity for IBM i and IBM Z hosts.

8.2 User and role management

During the planning phase (when you use the customization worksheet), list all users who need access to the DS GUI or DS Command-Line Interface (DS CLI). This action helps you manage secure authorization, which specifies the resource and access for different role-based users.

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. Assign one or more users to each of the following roles:

- ▶ The Administrator (admin) has access to several HMC or MC service functions and all storage image resources, except for specific encryption functions. This user authorizes the actions of the Security Administrator during the encryption deadlock prevention and resolution process.
- ▶ The Security Administrator (secadmin) has access to all encryption functions. A user with an Administrator role is required to confirm the actions that are taken by a user of this role during the encryption deadlock prevention and resolution process.
- ▶ The Physical operator (op_storage) has access to physical configuration service methods and resources, such as managing the storage complex, storage image, rank, array, and extent pool objects.
- ▶ The Logical operator (op_volume) has access to all service methods and resources that relate to logical volumes, hosts, host ports, logical subsystems (LSSs), and volume groups, excluding security methods.
- ▶ The Monitor role has access to all read-only, nonsecurity MC service methods, such as the **list** and **show** commands.
- ▶ The IBM Service role (ibm_service) has access to all MC service methods and resources, such as running code loads and retrieving problem logs. This group also has the privileges of the Monitor group, excluding security methods.
- ▶ The IBM Engineering role (ibm_engineering) has all access that the ibm_service group has plus more permissions to manage Fibre Channel (FC) Port settings, manage data-at-rest encryption, and modify Easy Tier settings.
- ▶ The Copy Services (CS) operator (op_copy_services) has access to all CS methods and resources, and the privileges of the Monitor group, excluding security methods.
- ▶ The Logical and Copy operator (op_volume and op_copy_services) has the combined access of the Logical operator and the Copy operator.

Important: Resource groups offer an enhanced security capability that supports the hosting of multiple customers with CS requirements. It also supports a single client with requirements to isolate the data of multiple operating system (OS) environments. For more information, see *IBM DS8000 Copy Services: Updated for IBM DS8000 Release 9.1*, SG24-8367.

For more information about the capabilities of certain user roles, see [User roles](#) or use the DS GUI Help function.

The DS8900F provides a storage administrator with the ability to create custom user roles with a fully customized set of permissions by using the DS GUI or DS CLI. This set of permission helps to ensure that the authorization level of each user on the system exactly matches their job role in the company so that the security of the system is more robust against internal attacks or mistakes.

You can also consider using a Lightweight Directory Access Protocol (LDAP) server for authenticating IBM DS8000 users. You can now take advantage of the IBM Copy Services Manager and its LDAP client that comes preinstalled on the DS8900F HMC. For more information about remote authentication and LDAP for the DS8900F, see *LDAP Authentication for IBM Storage DS8000 Systems: Updated for DS8000 Release 9.3.2*, REDP-5460.

8.3 Encryption

More planning is required if you intend to activate data-at-rest encryption for the DS8900F system. It is important to plan and configure data-at-rest encryption before you perform the logical configuration.

The following options are available:

- ▶ **Data-at-rest encryption:** DS8900F support for data-at-rest encryption consists of hardware-level, self-encrypting drives. The drive-based encryption is combined with an external enterprise-scale key management infrastructure to provide increased data security. The external key manager generates, protects, stores, and maintains encryption keys that are needed to encrypt information.

Since Release 9.2, the DS8900F system also offers data-at-rest encryption without needing external key servers. This feature is known as the *Local Key Management* feature. With the Local Key Management feature, the encryption keys are generated as local key group, and protected, stored, and maintained by the DS8900F. The Local Key Management feature does not require a specific license, but the option to use this feature must be selected during the initial order process of the DS8000 system. This feature cannot be activated by using a license function later. However, you can still decide to use external key server-based encryption. Each method is exclusive.

Additionally, DS8900F encryption also offers a simple, cost-effective solution for securely erasing any flash drive that is being retired or repurposed (cryptographic erasure). Full-disk-encryption drives are standard on the DS8900F system.

- ▶ **IBM Fibre Channel Endpoint Security:** Use this feature to encrypt data as it is transmitted between your IBM Z server and DS8900F. This encryption mechanism also uses external key servers. It can be enabled at any time.
- ▶ **Transparent Cloud Tiering (TCT):** The DS8900F system also provides support to encrypt data being transferred to the cloud when using the TCT function. TCT encryption can be enabled at any time, and relies on external key managers.

For more information, including considerations and best practices for DS8900F encryption, see 5.3.6, “Key manager servers for encryption” on page 161 and *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security (DS8000 Release 9.2)*, REDP-4500.

For more information about encryption license considerations, see “Encryption activation review planning” on page 162.

8.4 Network security

The security of the network that is used to communicate to and manage the DS8900F (specifically the HMC) is important depending on the client requirements. The DS8900F supports compliance with the National Institute of Standards and Technology (NIST) SP 800-131a standards.

Two components are required to provide full network protection:

- ▶ The first component is Internet Protocol Security (IPsec), and for Gen 2 security, IPsec v3 is required. IPsec protects network communication at the internet layer or the packets that are sent over the network. This configuration ensures that a valid workstation or server communicates with the HMC and that the communication between them cannot be intercepted.
- ▶ The second component is Transport Layer Security (TLS) 1.2, which provides protection at the application layer to ensure that valid software (external to the HMC or client) is communicating with the software (server) in the HMC.

Note: The details for implementing and managing security requirements are provided in *IBM DS8870 and NIST SP 800-131a Compliance*, REDP-5069.

8.5 Configuration flow

This section shows the list of tasks to perform when storage is configured in the DS8900F. Depending on the environment and requirements, not all tasks might be necessary.

Logical configuration can be performed by using the DS GUI, DS CLI, or a combination of both. Depending on your preference and experience, one method might be more efficient than the other. The DS8900F GUI provides a powerful yet simple process for logical configuration. If you use the DS GUI, not all of the steps that are listed in this book are explicitly performed by the user. For more information about the DS GUI, see Chapter 9, “IBM DS8900F Storage Management GUI” on page 233.

If you perform logical configuration by using the DS CLI, the following steps provide a high-level overview of the configuration flow. For more detailed information about using and performing logical configuration with the DS CLI, see Chapter 10, “IBM DS8900F Storage Management Command-line Interface” on page 339.

Here is the general configuration flow:

1. Install license keys: Activate the license keys for the DS8900F storage system. For more information about activating licensed functions, see Chapter 7, “IBM DS8900F features and licensed functions” on page 199.

Important: If data-at-rest encryption will be activated, the encryption configuration must be performed before starting the logical configuration.

2. Create arrays: Configure the installed flash drives as redundant array of independent disks (RAID) 6, which is the default and preferred RAID configuration for the DS8900F.
3. Create ranks: Assign each array as a Fixed-Block (FB) rank or a Count Key Data (CKD) rank.
4. Create extent pools: Define extent pools, associate each one with Server 0 or Server 1, and assign at least one rank to each extent pool. To take advantage of storage pool striping, you must assign multiple ranks to an extent pool.

Important: If you plan to use IBM Easy Tier (in particular, in automatic mode), select the **All pools** option to receive all of the benefits of Easy Tier data management.

5. Consider other controls and monitoring when working with space-efficient volumes. For more information, see *IBM DS8880 Thin Provisioning (Updated for Release 8.5)*, REDP-5343.
6. Configure the FC ports: Define the topology of the FC ports. The port type can be Fibre Channel Protocol (FCP) or Fibre Channel connection (IBM FICON).
7. Create the volume groups for open systems: Create volume groups where FB volumes are assigned.
8. Create the host connections for open systems: Define open systems hosts and their FC host bus adapter (HBA) worldwide port names (WWPNs). Assign volume groups to the host connections.
9. Create the open systems volumes: Create striped open systems FB volumes and assign them to one or more volume groups.
10. Create the IBM Z logical control units (LCUs): Define their type and other attributes, such as subsystem identifiers (SSIDs).
11. Create the striped IBM Z volumes: Create IBM Z CKD base volumes and parallel access volumes (PAV) aliases for them.

8.6 General storage configuration guidelines

Observe the following general guidelines when storage is configured in the DS8000:

- ▶ To achieve a well-balanced load distribution, use at least two extent pools (also known as a *pool pair*), each assigned to a different internal server. If CKD and FB volumes are required on the same storage system, configure at least four extent pools: Two for FB and two for CKD.
- ▶ The volume type for the first volume that is created in an address group is either FB or CKD. That volume type determines the type for all other volumes (FB or CKD) in the entire address group. A volume is one of 256 in an LSS or LCU. An LSS is one of 16 in an address group (except address group F, which has only 15 LSSs). For more information about LSSs and address groups, see 4.4.5, “Logical subsystems” on page 130.
- ▶ Volumes of one LCU or LSS can be allocated on multiple extent pools in the same rank group.
- ▶ Assign multiple ranks to extent pools to take advantage of storage pool striping. Additionally, assign ranks from multiple device adapter (DA) pairs to an extent pool to spread the workload and increase performance.
- ▶ The following options are available for FB pools:
 - Create a volume group for each server unless logical unit number (LUN) sharing is required.
 - Assign the volume group for one server to all of its host connections.
 - If LUN sharing is required, the following items provide two possible use cases (Figure 8-1 on page 231):
 - Application sharing: Create one volume group for each server. Assign the shared volumes in each volume group. Assign the individual volume groups to the corresponding server’s host connections. The advantage of this option is that you can assign private and shared volumes to each host.
 - Clustering: Create one common volume group for all servers. Place the shared volumes in the volume group and assign the volume group to the host connections.




Figure 8-1 LUN configuration for shared access

- ▶ The following options are available for FC ports:
 - Configure a port to be FICON or FCP.
 - Distribute host connections of each type (FICON and FCP) evenly across the I/O enclosures.
 - Ensure that each host is connected to at least two different host adapters in two different I/O enclosures for redundancy and availability.
 - As a best practice, allow hosts to access all host adapter ports of the system.

Note: Avoid intermixing host I/O with CS I/O on the same ports for performance reasons.



IBM DS8900F Storage Management GUI

This chapter describes the DS8000 Storage Management DS GUI that is available with the IBM DS8900F system. The DS GUI enables you to perform most functions on the storage system, including:

- ▶ Initial system setup for a new installation
- ▶ Activation of licensed functions
- ▶ Simplified and advanced logical configuration for open systems and IBM Z
- ▶ Graphical view of system resource availability
- ▶ System status and events viewer
- ▶ Access to advanced help and IBM Documentation
- ▶ Performance monitoring dashboard and statistics offload
- ▶ Custom user roles for more robust security of the system

This chapter covers the following topics:

- ▶ Introduction
- ▶ DS GUI: Getting started
- ▶ System configuration overview
- ▶ Logical configuration overview
- ▶ Logical configuration for open systems volumes
- ▶ Logical configuration for Count Key Data volumes
- ▶ Expanding volumes
- ▶ Deleting a pool
- ▶ Deleting volumes
- ▶ Reinitializing a thin-provisioned volume
- ▶ Easy Tier support
- ▶ Monitoring system health
- ▶ Performance monitoring
- ▶ Fibre Channel error rate statistics
- ▶ Providing feedback

9.1 Introduction

The DS GUI is packed with usability features that offer an easy-to-use interface for administration of the DS8900F.

The DS GUI was designed and developed with three major objectives:

- ▶ **Speed:** A graphical interface that is fast and responsive.
- ▶ **Simplicity:** A simplified and intuitive design that can drastically reduce the time that is required to perform functions with the system, which reduces the total cost of ownership (TCO).
- ▶ **Commonality:** Use of common graphics, widgets, terminology, and metaphors that facilitate the management of multiple IBM storage products and software products. The DS GUI that was introduced with Release 9.0 enables a consistent graphical experience and easier switching between other products like IBM FlashSystem®, IBM Spectrum Virtualize, or IBM Spectrum Control and IBM Storage Insights.

Based on these objectives, following the initial setup of the storage system, a system administrator can use the DS GUI to complete the logical configuration and then prepare the system for I/O. After the initial setup is complete, the system administrator can perform routine management and maintenance tasks with minimal effort, including the monitoring of performance, capacity, and other internal functions.

Logical storage configuration is streamlined in the DS GUI for ease of use. The conceptual approach of *array site*, *array*, and *ranks* is streamlined into a single resource, which is referred to as an *array* (or *managed array*). The storage system automatically manages flash adapter pairs and balances arrays and spares across the two processor nodes without user intervention.

Creating usable storage volumes for your hosts is equally simplified in the DS GUI. The system can automatically balance volume capacity over a pool pair. If custom options are required for your workload, the DS GUI can override defaults and customize your workload needs.

Configuring connections to hosts is also easy. Host ports are updated automatically and host mapping is allowed at volume creation.

The overall storage system status can be viewed at any time from the dashboard window. The dashboard presents a view of the overall system performance when a system administrator first logs in for a picture of the system status. This window also contains a “Hardware View” and a “System Health View” that displays the status and attributes of all hardware elements on the system.

Additionally, functions that include user access, licensed function activation, setup of encryption, IBM Fibre Channel Endpoint Security, and remote authentication, and modifying the power or Fibre Channel (FC) port protocols are available to the system administrator.

All functions that are performed in the DS GUI can be scripted by using the DS Command Life Interface (DS CLI), which is described in Chapter 10, “IBM DS8900F Storage Management Command-line Interface” on page 339.

9.2 DS GUI: Getting started

This section describes how to accomplish the following tasks:

- ▶ Accessing the DS GUI
- ▶ System Setup wizard
- ▶ Configuring Fibre Channel port protocols and topologies
- ▶ Managing and monitoring the storage system
- ▶ Storage Monitoring and Servicing from the Unified Service GUI

9.2.1 Accessing the DS GUI

The DS8000 Storage Management GUI (DS GUI) is a web-based GUI that is installed on the Hardware Management Console (HMC). You can access the DS GUI from any network-attached computer by using a supported web browser. Examples of the minimum levels of the supported web browsers at the time of writing are:

- ▶ Mozilla Firefox 68
- ▶ Microsoft Edge 44
- ▶ Google Chrome 76

For any specific requirements on your browser, see [IBM Documentation](#).

On a new storage system, the user must log on as the administrator. The password expires immediately, and the user is forced to change the password.

The initial view of the system is shown in Figure 9-1.

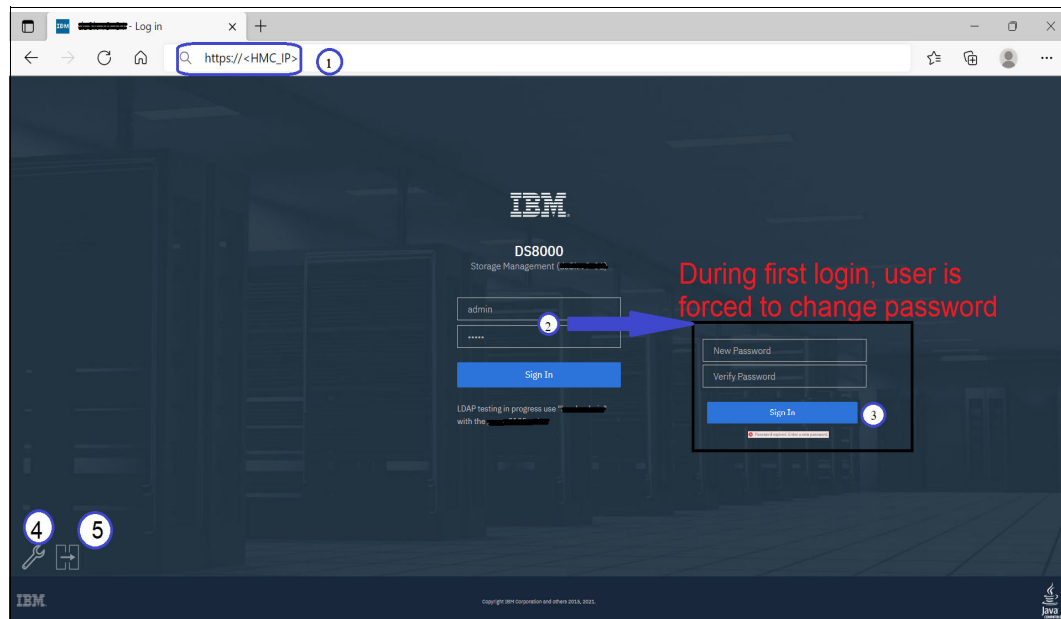


Figure 9-1 Storage Management GUI initial login

Figure 9-1 on page 235 shows the following components:

1. DS GUI URL:

https://<HMC_IP>

This DS GUI URL is configured by the IBM Systems Service Representative (IBM SSR) during system installation.

2. Initial log in to the new system:

- User = admin
- Password = admin

3. The default password expires immediately, and the user is forced to change the password for the admin user ID.

4. IBM Support personnel use the Service Management Console (wrench) icon to log in to the system as needed.

5. Clicking the IBM Copy Services Manager icon opens a window to the IBM Copy Services Manager Console that is installed on the HMC, as shown in Figure 9-2.

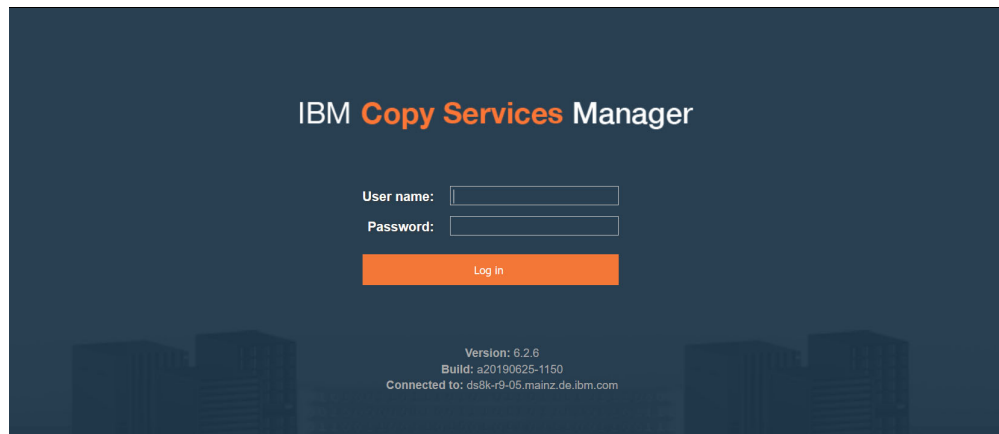


Figure 9-2 IBM Copy Services Manager window started from the DS GUI

9.2.2 System Setup wizard

For a new storage system installation, the DS GUI starts the *System Setup wizard*. The wizard starts automatically after the admin password is changed and a user with the Administrator role and authority logs in.

The wizard guides the admin user through the following tasks:

1. Set the system name.
2. Activate the licensed functions.
3. Provide a summary of actions.

The System Setup window opens with the Welcome pane, as shown in Figure 9-3.

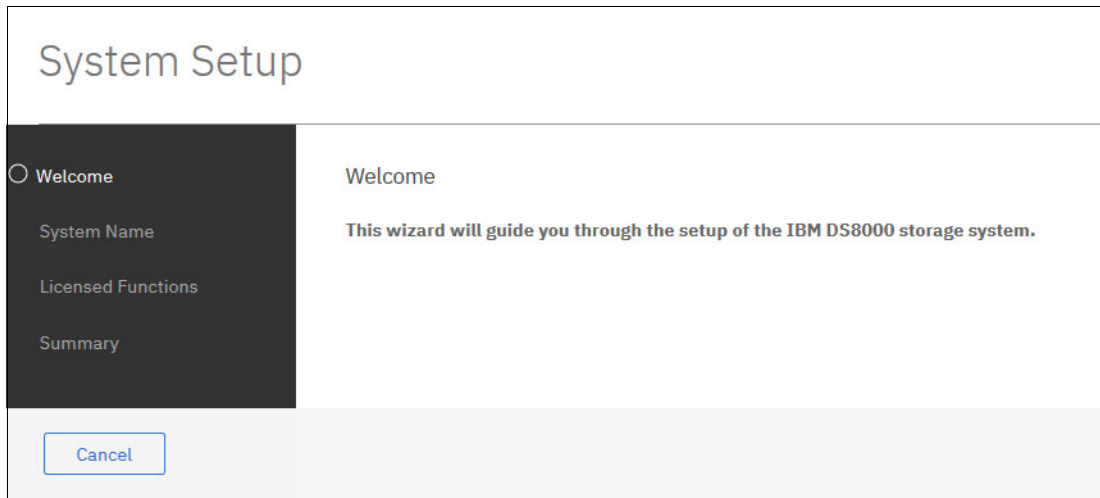


Figure 9-3 System Setup: Welcome pane

Starting with the Welcome pane, complete the following steps:

1. Click **Next**. The System Name window opens. The default entry is the storage system serial number, which is shown in Figure 9-4. The user can create a preferred system name that complies with the name convention that is used in the environment.

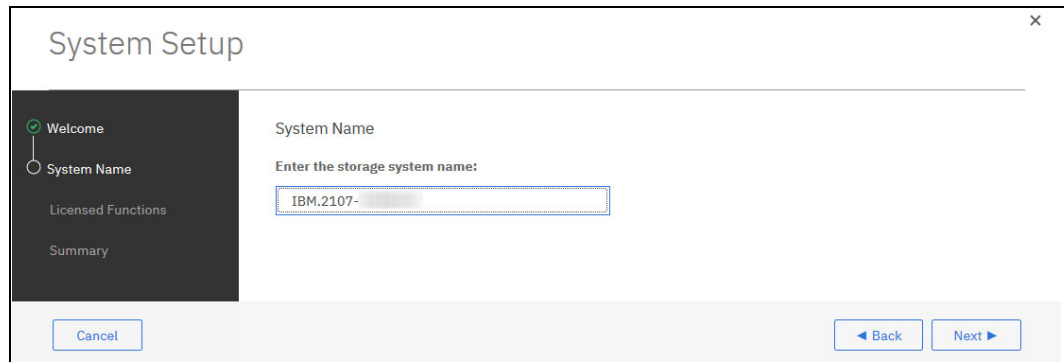


Figure 9-4 System Setup: System Name window

2. Click **Next**. The Licensed Functions window opens. Click **Activate Licensed Functions**.
3. The Activate Licensed Functions window opens. Keys for licensed functions that are purchased for this storage system must be retrieved by using the Machine Type, Serial Number, and Machine Signature. The keys can be stored in a flat file or an XML file. Licensed function keys are downloaded from the [data storage feature activation \(DSFA\) website](#).

Figure 9-5 shows the Activate License Functions window, the help information, and the activation of the licensed functions. You can import the keys by selecting the **Folder** icon or paste the keys in the window. For more information about licensed functions, see Chapter 7, “IBM DS8900F features and licensed functions” on page 199.

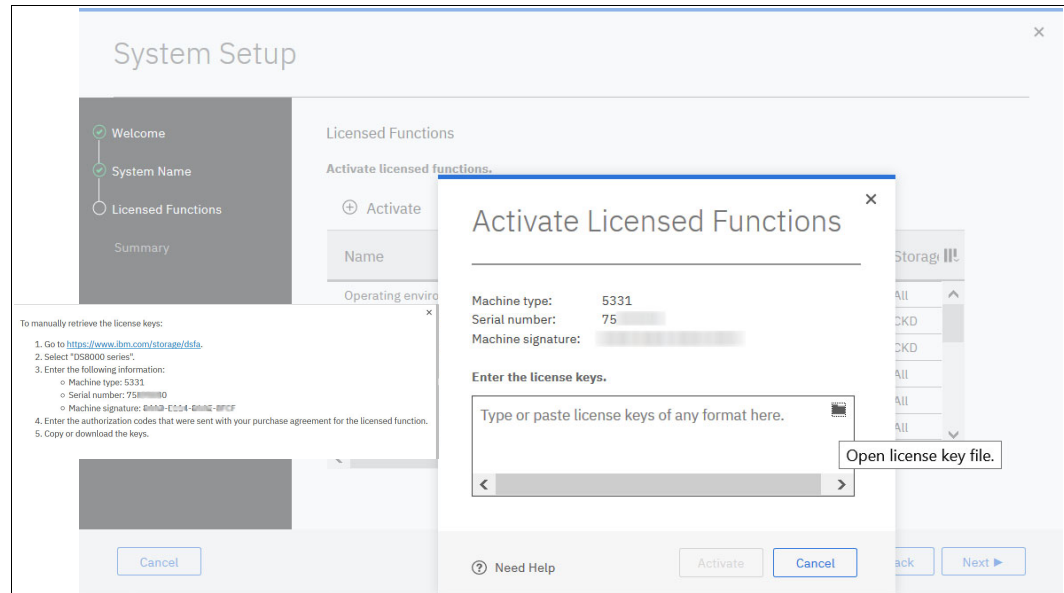


Figure 9-5 System Setup: Activate Licensed Functions

4. When the license keys are entered, click **Activate** to enable the functions, as illustrated in Figure 9-6.

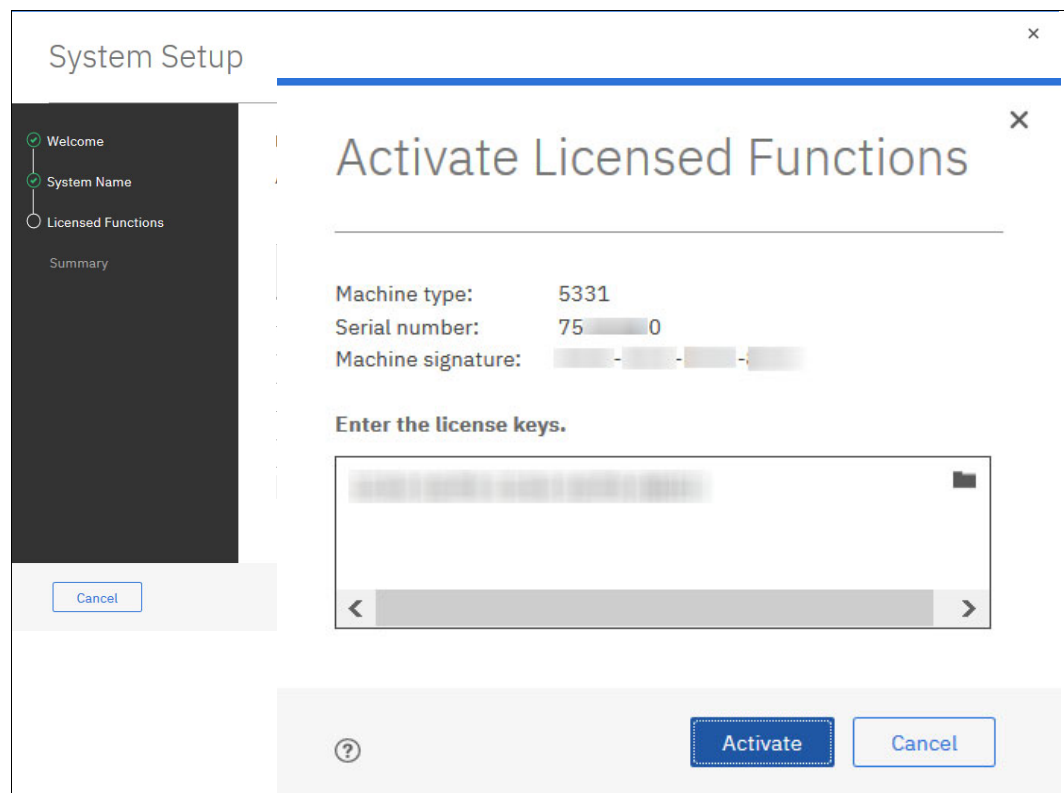


Figure 9-6 System Setup: Enabling Licensed Functions

- Click **Next** to open the Summary window, which is shown in Figure 9-7. If everything looks correct, click **Finish** to exit the System Setup wizard. After the wizard is closed, the System window opens.

Note: The Summary shows licenses for basic functions. The list might include some advanced functions such as Copy Services (CS), Z Synergy Services (zsS), and IBM Copy Services Manager on HMC if the corresponding licenses were activated.

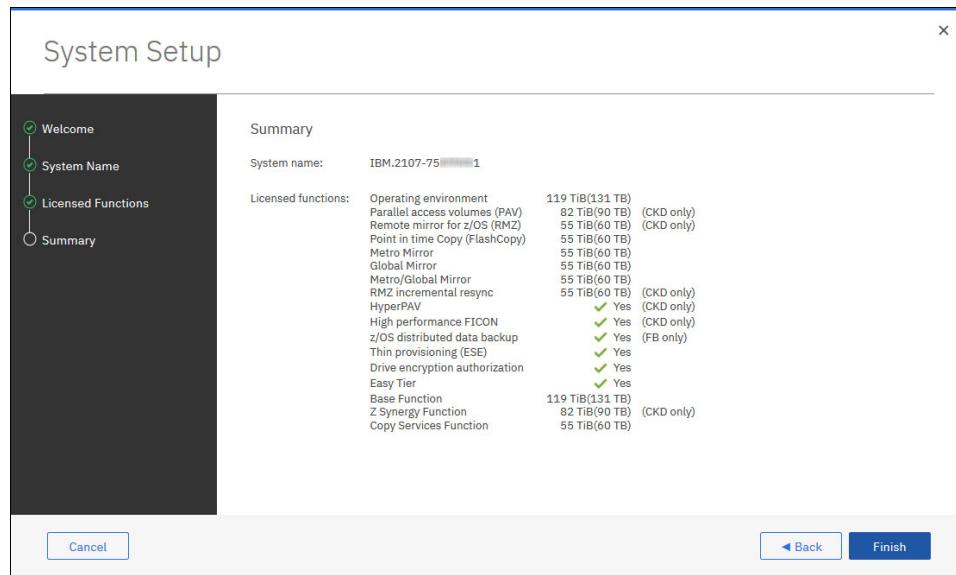


Figure 9-7 System Setup: Summary

9.2.3 Configuring Fibre Channel port protocols and topologies

After the initial setup is complete, you can configure the FC port topologies for your host attachment requirements (open system or IBM Z).

There are several ways to do this task in the DS GUI:

- ▶ A quick way to perform this action is by clicking **Actions** on the Dashboard window and selecting **Modify Fibre Channel Port Protocols**, as shown in Figure 9-8.

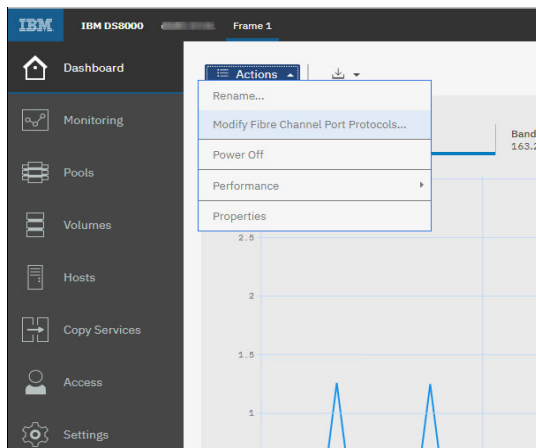


Figure 9-8 Quick way to modify the fiber adapter port protocols

- ▶ Another option is to click **Settings** → **Network** → **Fibre Channel Ports**, select a specific port from the list, and select **Actions** to set the ports, as shown in Figure 9-9.

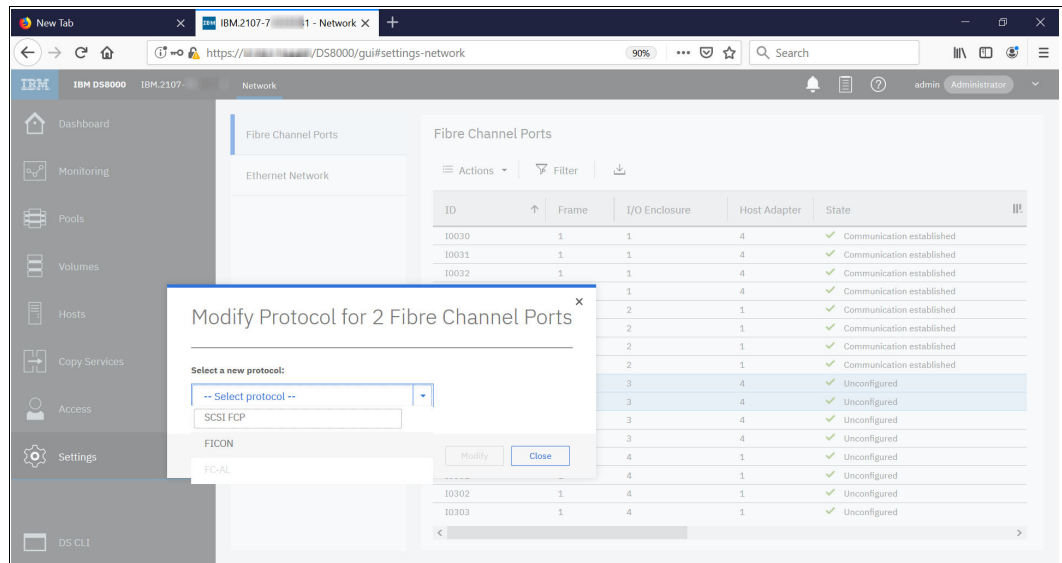


Figure 9-9 Configuring the Fibre Channel port topology

- ▶ You can also configure the port topology from the System view or from the System Health overview, as shown in Figure 9-10.




Figure 9-10 Modifying the FC port protocol from the System Health overview

You can also perform this configuration during logical configuration. For more information, see 9.6.4, “Creating FB host attachments” on page 282, and 9.13, “Monitoring system health” on page 309.

9.3 Managing and monitoring the storage system

When the initial setup of the system is complete, the Dashboard opens. This view is what opens after a user logs in. It displays the main hardware components of the DS8900F system and shows the status of the system hardware, a view of the overall system performance for a quick picture of the system performance, and a summary of the system capacity. You can return to this window at any time by clicking the **Dashboard** icon in the upper left of the GUI.

Note: Different users might have a limited view of the Dashboard when logging in, depending on their role. Most of the material that is documented here describes what the Administrator role sees.

From the DS GUI, the administrator can manage the system by performing actions for various activities, such as:

- ▶ Logical configuration
- ▶ Controlling how the system is powered on or off
- ▶ Modifying the FC port protocols and customer network settings
- ▶ Modifying Easy Tier settings
- ▶ Viewing system properties
- ▶ Displaying performance monitoring graphs
- ▶ Accessing the embedded DS CLI
- ▶ Viewing the status and properties of some CS functions, such as FlashCopy and mirroring

Figure 9-11 presents a high-level overview of the System window and all the objects that can be accessed from this window.

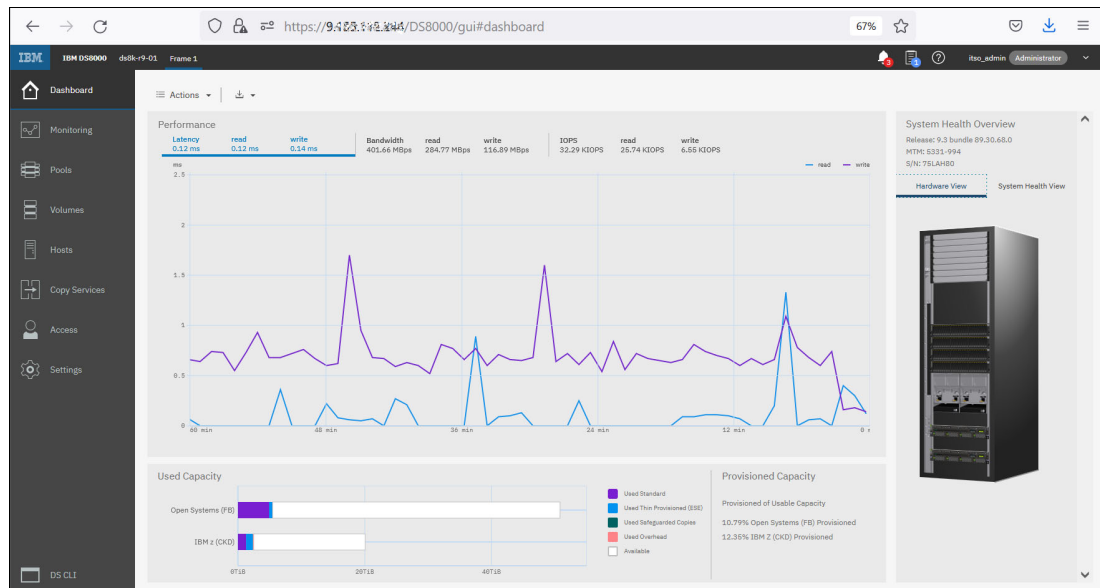


Figure 9-11 Initial System view

Note: The menu items and actions that are shown in the DS GUI depend on the role of the user that is logged in, and they can vary for each user. For more information about user roles, click **Help** at the upper right of the DS GUI window and search for *user role*.

This initial view of the system provides access to a wealth of information in a quick and easy to use view. Items that are included in this initial view are:

► **Dashboard icon:**

- Click the **Dashboard** icon from anywhere in the DS GUI to return to the system Dashboard. The Dashboard provides an overall view of the DS8000 system.
- At the top, the serial number of the DS8000 system you are working on is shown.
- Actions menu:

- **Rename.**

Change the name of the DS8900F storage system.

- **Modify Fibre Channel Port Protocols.**

Select the protocol that is used by FC ports to connect the storage system to a host or to another storage system. The user can also display the properties of the selected port, which opens the same view that you get when you select **Settings** → **Network** → **Fibre Channel Ports**.

- **Power Off/On.**

Initiate a power-off or power-on of the DS8000 storage system.

- **Performance**

Shows graphs that display performance metrics for I/O operations per second (IOPS), Latency, Bandwidth, and Caching.

- **Properties.**

This view displays the system properties that are shown in Table 9-1.

Table 9-1 System properties

System property
System name
Current state
Product type
Release and bundle version of the IBM system software
Machine type and model (MTM)
Unique serial number
Machine Signature (string of characters that identifies a DS8000 storage system), which is used to retrieve license keys from the DSFA website
Worldwide node name (WWNN) of the system
Hardware component summary (such as processor type, total subsystem memory, raw data storage capacity, or number of FC ports)
Last power-on timestamp

- Export reports.

Since Release 9.2, every user role, including Monitor, can download reports such as the System Summary comma-separated values (CSV) report, Performance Summary CSV report, Easy Tier Summary report, and FC Connectivity report. Previously, exporting reports was limited to the Storage Administrator role.

To select multiple options that will be saved in a compressed CSV file, click the **Download** icon. The options include the System Summary, Performance Summary, Easy Tier Summary, and the FC Connectivity report.

Note: The System Capacity section in the System Summary CSV is composed of consolidated data. The System Capacity, Used IBM Z (Count Key Data (CKD)) Capacity, and Used Open Systems (Fixed-Block (FB)) Capacity sections in the System Summary CSV are now combined into one section that is called System Capacity. All sections are now shown with the column headers listed even if there are no items in the list. For example, the FlashCopy section is shown even if no FlashCopy relationships are present on the system.

Another enhancement since Release 9.2 is the FC connectivity report, which is a compressed file that contains a report that shows one row for every connected path between the DS8000 and a host system, a switch, or another DS8000. It also shows the status of these paths and their security settings.

- Alerts.

Highlights events at which you should look with priority.

- Suggested tasks.

If there are any suggested tasks, they are indicated here.

- Help icon.

Clicking this icon opens a drop-down menu that includes the following options:

- The first option in this menu is context-sensitive, and it depends on the window in which you are currently in the DS GUI. Clicking this option opens a separate window to the DS GUI IBM Documentation for that specific window, and describes resources and actions for that specific window.
- What's New: This option provides descriptions of new features that are in the installed release of the GUI with a link to the DS GUI IBM Documentation for each of the new functions. Also provided is a link to learn more about the newest functions.
- Help Contents: This option opens a separate window for the DS GUI IBM Documentation main window.
- About DS8000: This option opens the DS8000 Storage Management window, which shows information that is specific to the system, including code level, product, and machine serial number.

For more information about the DS8900F help functions, see 9.3.1, “Storage Monitoring and Servicing from the Unified Service GUI” on page 248.

- Hardware View.

A graphic of the DS8900F frame is shown.

Hover your cursor over an enclosure, node, HMC, or the processor nodes, and a Properties window opens that shows basic information about that particular component in the frame.

If a magnifying glass icon appears when hovering over a component, click it to bring the component to the center of the window. You can now hover your cursor over the subcomponents to see a Properties window that shows basic information about it. Click anywhere on the window to close the view.

When viewing an FC port on a host adapter within an enclosure, you can right-click a port to modify the FC port protocol and view properties and events.

The following information is typically shown for an enclosure:

- Enclosure.
- ID.
- State, for example, online, offline, or service required.
- Enclosure-specific data.
- Enclosure serial number.

The following information is typically shown for a node:

- ID.
- State, for example, online, offline, and service required.
- DS8900F release.
- Data that is specific to an IBM Power server.

For more information about the Hardware view, see 9.13, “Monitoring system health” on page 309.

– Performance.

This graph displays the performance of the system reads/writes based on latency, bandwidth, or IOPS. Click one of the sections to view the graph.

– System Health Overview.

Clicking this tab opens the System Health Overview in a table format. In addition to displaying the machine type, model, serial number, and microcode level that is installed on the system, this section also shows the overall health status and quantity of various hardware components in the system, including processor nodes, HMCs, storage enclosures, drives, I/O enclosures, host adapters, device adapters (DAs), and FC ports. Clicking each component shows the details for that component in a table. Click **Close** to exit.

– Capacity.

The capacity area consists of a bar graph displaying the used capacity of the system by type (Open Systems or IBM Z) and unassigned capacity. Next to the bar graph is a legend for each of the colors that are used on the graph. On the right side of this area is the provisioned capacity summary of the system.

► Monitoring menu:

- Events: This option opens the Events window, which displays all of the events in severity order.
- Performance: This option opens the Performance (statistics) window.

- ▶ Pools menu:
 - Arrays by Pool: Access this view to see all the pools on the system along with the arrays that they contain. Use this view to access actions that can be performed on pools and arrays. This view shows any unconfigured arrays.
 - Volumes by Pool: Access this view to see all the pools on the system along with the volumes that they contain. Use this view to access actions that can be performed on pools and volumes.
- ▶ Volumes menu:
 - Volumes: Access this view to see all the volumes on the system. Use this view to access all actions that can be performed on volumes, such as create, modify, or delete volumes.
 - Volumes by Pool: This view is the same one that is described in the Pools menu.
 - Volumes by Host: Access this view to see volumes that are based on the host or host cluster to which they are assigned and all volumes that are not mapped to a host. Use this view to access all actions that can be performed on volumes and hosts or host clusters.
 - Volumes by LSS: Access this view to see volumes that are based on the logical subsystem (LSS) to which they belong. Use this view to access all actions that can be performed on volumes and LSSs.
- ▶ Hosts menu:
 - Hosts: Access this view to see all the hosts and host clusters on the system. Use this view to access all actions that can be performed on hosts and host clusters, such as create, modify, or delete hosts or host clusters, and the state of host port logins.
 - Volumes by Host: The same view that is described in the Volumes menu.
- ▶ Copy Services menu:
 - FlashCopy: The FlashCopy window provides details about FlashCopy relationships.
 - Mirroring: The Mirroring window provides details and status information about Remote Mirror and Remote Copy volume relationships.
 - Mirroring Paths: The Mirroring Paths window displays a list of existing Remote Mirror and Remote Copy path definitions.
- ▶ Access menu:
 - Users:

Only users with the administrator role can access this menu. This menu opens the Users window. A system administrator can use this menu to perform the following actions:

 - Create user accounts.
 - Set a user account role.
 - Set temporary passwords (to be reset at first use by the new account).
 - Modify an existing user account role.
 - Reset an existing user account password.
 - Disconnect a user account.
 - Determine a user account connection (DS CLI or GUI).
 - Remove user accounts.

– Roles:

A storage or security administrator can set up user roles in the GUI or CLI with a fully customizable set of permissions to ensure that the authorization level of each user account matches their job role to ensure that the security of the system is more robust against internal attacks or mistakes. The following actions can be taken against roles:

- Create custom user roles.
- Modify remote mappings.
- Delete roles.

Restriction: Default roles cannot be deleted.

- View permissions for each role.
- View properties for each role.

– Remote Authentication.

This menu allows the user to set up remote authentication through a central repository.

► Settings menu:

– Network:

- Modify the FC ports protocol for a selected port or group of ports.
- Display error rates for a selected port or group of ports.
- Display a single port's properties.
- View the current Ethernet network information and change settings for both HMCs.

– Security:

- Manage encryption for the storage system, that is, data at rest encryption and IBM Fibre Channel Endpoint Security.
- Manage local password rules (such as password minimum length, expiration, and age).
- Manage the communication certificate on the HMC to enable HTTPS connections with the HMC.

– System:

- Licensed Function.

This window shows a summary of the activated licensed functions. When you click the **Activate** option, the Activate Licensed Functions window opens, where you can activate more licenses. For a full description about activating licensed functions, see Chapter 7, “IBM DS8900F features and licensed functions” on page 199.

- Easy Tier.

You can use this function to enable and configure Easy Tier to improve performance by managing or monitoring the volume capacity placement in pools. The Easy Tier Heat Map Transfer Utility (HMTU) can also be enabled.

- zHyperLink.

You can use this function to enable or disable the zHyperLink as I/O Read Enabled and I/O Write.

- Date and Time.

Set the time zone and set the date and time manually for the system or enter the IP address of a Network Time Protocol (NTP) server.

- **Advanced.**
On the **Advanced** tab of the System settings window, you can allow service access, enable ESSNet CS, set the IBM i serial number prefix, enable control-unit initiated reconfiguration (CUIR) for IBM Z, and select the power control mode for the storage system. In addition, you can manage service settings and work with other settings for your system.
- **Notifications:**
 - **Call Home.**
You can enable Call Home on your Management Console (MC) to send an electronic Call Home record to IBM Support when there is a problem within the storage complex.
 - **Syslog.**
You can define, modify, or remove syslog servers. You can also enable extra security with Transport Layer Security (TLS) for the syslog.
- **Support:**
 - **IBM Remote Support Center (RSC).**
You can configure RSC to allow IBM Support to remotely access this system to quickly resolve any issues that you might be having. You can choose that the RSC connection stays open always, close 2 hours after IBM support is logged off, or closed. For added security, you can require an access code for remote support.
 - **Assist On-site.**
You can configure the Assist On-site (AOS) feature, which allows IBM Support to remotely access the MC and storage system. Choose an option to stop, start, or restart the AOS service.
 - **Troubleshooting.**
You can restart the local or remote HMC to correct communication issues between the storage system and HMC.

You can refresh the GUI cache if the data in the Storage Management GUI is not in sync with data in the DS CLI or IBM Spectrum Control.

You can restart the web servers and communication paths that are used by IBM Enterprise Storage Server Network Interface (IBM ESSNI).
- **GUI Preferences:**
 - **Login Message.**
You can enter a message that is displayed when users log in to either the DS GUI or an interactive DS CLI session.
 - **General.**
You can set the default logout timeout and chose to show suggested tasks.

► Embedded DS CLI.

Click the **Embedded DS CLI** icon to open a DS CLI session from within the DS GUI, as shown in Figure 9-12. The version of embedded DS CLI is the most current one that is available for the microcode that is installed on the DS8900F.

The DS CLI commands can be run conveniently from the GUI with the least amount of response time by avoiding network hops when using remote DS CLI.

DS CLI scripts can also be run from the embedded DS CLI on the DS GUI. The script must be text-based and on the workstation running the DS GUI session.

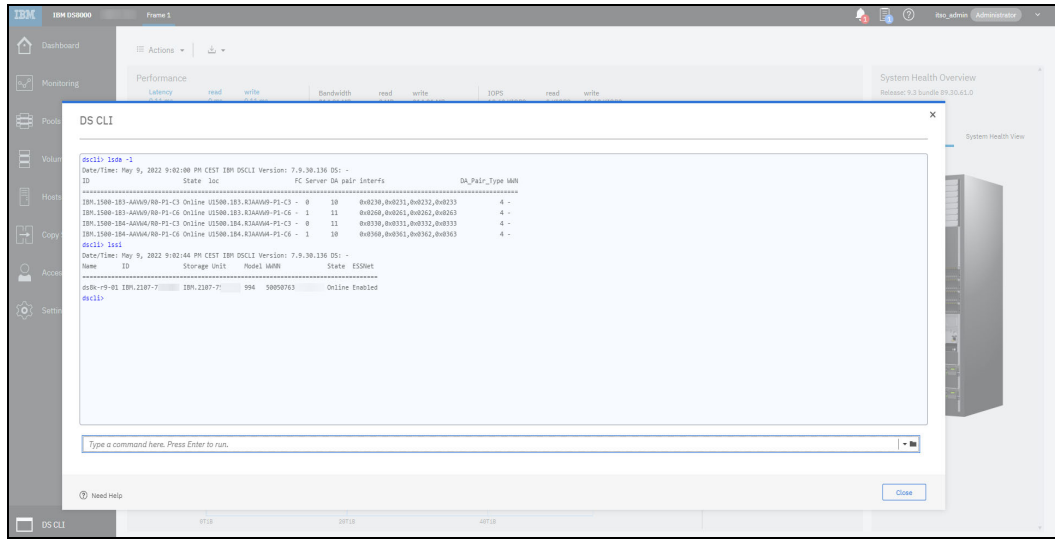


Figure 9-12 Embedded DS CLI

9.3.1 Storage Monitoring and Servicing from the Unified Service GUI

The Unified Service GUI, shown in Figure 9-13 on page 249, provides access to all service functions and tools. Before Release 9.2, all these features were available for use only through the Service Web User Interface (WUI) on the HMC. All these functions help IBM Support representatives perform specific tasks like data collection, event management, miscellaneous equipment specification (MES), model conversions, and microcode and hardware upgrades.

Important: This Unified Service GUI should be accessed by IBM Support representatives or under the supervision of IBM support representatives only.

Service Dashboard was a recent add-on to the DS8900F GUI after Release 9.2. Service Dashboard is accessible by an *IBM service role privilege user only* on the DS8900F GUI.

All functions in the service dashboard are added from the WUI. The categorization and functions of each attribute under the service dashboard perform functions like under the WUI, which provides ease of management and more privilege to customers. These functions are still available through WUI and HMC access, but they are added to the DS8900F GUI interface too.

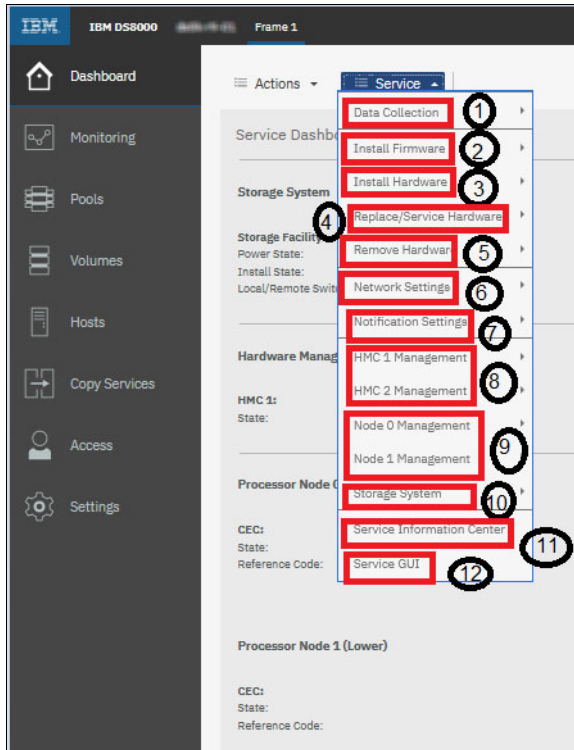


Figure 9-13 Unified Service GUI

1. Data collection:

- a. Perform Data collection on Demand: With this option, you can generate and offload the existing complete PE package with different formats.
 - i. General PE Package: Collects data for any service actions, such as installation or removal, MES, repair, and code load. This package contains the data of reliability, availability, and serviceability (RAS) and functional code components (not including state saves).
 - ii. Client User Interface (ESSNI, DS GUI, or API) Package: Collects data for problems that originate from using customer user interface applications, such as DS GUI. This package contains the data of various software components running on HMC, such as DS GUI, ESSNI, and RESTful APIs. DS CLI traces are not on the HMC.
 - iii. All data packages for removable media offload only option: Collects both the PE packages and off loads the data to removable media.
- b. Offload Data by Area: With this option, you can download an area package of any of the following line items. This option is intended only as a substitute for the “PE Package/ Full Data collection On-Demand” function. Many of the listed items contain duplicate data. If you require more than one or two of these items, then it is a best practice to cancel this function and use the “PE Package /Full Data collection on-demand” function.
 - i. ESSNI.
 - ii. DEVICE Object Mismatch.
 - iii. Panics.
 - iv. Failed Service Action (Repair).
 - v. CDA.

- vi. Logical configuration Errors.
 - vii. Machine Down (HMC still up).
 - viii. FCIC Loop Issues.
 - ix. MES Power/Cache.
 - x. MES Storage Enclosure.
- c. Individual file offload: With this option, you can off load the selected file from the storage facility image. When prompted by IBM Support personnel, service users off load the required files from the different components of the storage facility image like from HMC or the LPARs.
 - d. Manage Dumps: With this option, you can generate and manage the individual dumps from the physical server (CPC).
 - e. Process state saves and OnDemand Data Dump (ODD) dumps:
 - i. Process LPAR State Saves: With this option, you can generate new state saves for the storage facility image or offload the existing state saves for troubleshooting purposes when prompted to do so by IBM Support personnel.
 - ii. Process LPAR ODD DUMPS: With this option, you can generate an ODD for the storage facility or off load the existing data dump for troubleshooting purposes when prompted to do so by IBM Support personnel.
 - iii. Process LPAR Lightweight ODD Dumps: With this option, you can generate lightweight data dumps for the LPARs, or a user can offload the existing lightweight dump for troubleshooting purposes when prompted to do so by IBM Support personnel.
 - f. Process Device Adapter and POST State Saves: With This option, you can generate state saves for the HA and DA adapter, or a user can offload the existing state saves for troubleshooting purposes when prompted to do so by IBM Support personnel.
 - g. Process Storage Enclosure State Saves: With this option, you can generate state saves for the selected storage enclosure with offloading the state saves data.
 - h. Process Drive State Saves: With this option, you can generate state saves for the selected disk drive module (DDM) by offloading the state saves data.
2. Install firmware: This option was available in the WUI console under the updates option. Under the Updates window, you can find the similar functions under the Install Firmware section in the DS8900F GUI service dashboard.
 - a. Verify Bundle Installed: Validates and lists the Active Firmware version of all components in the storage facility from the selected bundle's expected firmware.
 - b. Display Storage Facility Code Levels: Displays the Active Firmware version of all components in the storage facility.
 - c. Display Preload Status: Lists the status of any preloads.
 - d. Run CDA Preverify: Performs a full system scan to validate all things are working as expected. Generates the logs.
 - e. Acquire Code bundle: This function was available while selecting the HMC. With it, you can acquire the new release code that is available through the HMC by using a CD-ROM or an FTP server.
 - f. Enable/Disable File server: Enables and disables the file server capabilities on the HMC, which allows a Remote File Download Client request.
 - g. Remote file Download: Activates the file server capabilities on the designated HMC IP address.

- h. Update HMC code: Updates the HMC code from downloaded code in the HMC.
- i. Update Storage Facility Code: A service user can select the available facility code and either **Distribute only**, **Activate only**, or both.
- j. Select and Install corrective Services: With this option, you can install code add-ons for certain components.
- k. ICS Utilities.
- l. Prepare HMC Upgrade: Selects the recovery image for HMC for an upgrade activity. After a restart, the HMC begins the upgrade.
- m. Rebuild Peer HMC: Rebuilds the Peer HMC.
- n. Advanced Utilities:
 - i. Advanced configuration, Install Corrective Service, Display Library Contents, Clear Library Contents, Delete Release Bundle and Package, and Delete a Recovery Image.
 - ii. Display/Update Bootlist image & Reverse eServer™ Firmware is available in the WUI update section under the HMC in Backlevel utilities, and in the DS8900F GUI under Advanced Utilities.
 - iii. Display/Reset CDA SFI attributes & Reset Serviceable Event Tracking is available in the WUI update section under HMC in Miscellaneous utilities, and in the DS8900F GUI under Advanced Utilities.
- o. CCL Utilities: With concurrent code load, you can perform CCL I/O Enclosure, IBM Power firmware, and storage enclosure updates.
- p. Non-concurrent code load (NCCL) Utilities: You can perform NCCL for the following components:
 - i. NCCL SFI code Activation Single LPAR: No IML, NCCL SFI code Activation Single LPAR – Resume, and NCCL SFI code Activation Single LPAR – Start CPSS.
 - ii. NCCL eServer firmware update, NCCL eServer firmware Single Node Update, NCCL I/O Enclosure firmware update, and NCCL Power firmware update.
 - iii. NCCL SFI Code Activation: IML, and NCCL SFI Code Activation – NO IML.
- 3. Drive Utilities: You can perform Display Drive Code Levels, Display Drive Update Status, Run Drive Pre-verify, CCL Update Drive Code Level, NCCL Update Drive Code Level, and Terminate Drive Update on Drives.
- 4. Install Hardware: With this option, you can use the hardware component installation assistance wizard:
 - a. Storage Facility Field Install.
 - b. Generate Install Report: This section is available under the Storage facility Management section in the WUI interface.
 - c. View/Certify Drive: View and certify installed drives.
 - d. You can install the option Open wizard to help with the installation or MES upgrade of the following items:
 - i. Install I/O Enclosure or Components.
 - ii. Install Rack Power Components.
 - iii. Install Storage Enclosure or Drives.
 - iv. Install Expansion Rack.
 - e. Storage Facility Conversion: With this option, you can do a model conversion.

5. Replace service Hardware: Replace or service hardware components like I/O enclosures, storage enclosures, CPCs, and racks:
 - a. Manage Serviceable Events: View details or close open events. Available under the Service management section of the WUI.
 - b. Exchange CEC Components: A wizard helps to exchange the CPC component.
 - c. Exchange I/O Enclosures and Components: A wizard helps to exchange the I/O enclosures and components.
 - d. Exchange Rack Components: A wizard helps to exchange the rack components.
 - e. Exchange Storage Enclosures and Components: A wizard helps to exchange storage enclosures.
 - f. Manage field-replaceable units (FRUs).
 - g. Activate/Deactivate Resources: Power supply and Remote Procedure Call (RPC).
6. Remove Hardware: This feature should be used during an RPQ or next-level support supervision. These wizards help to remove the following listed devices:
 - a. Remove CPC Enclosure or Components.
 - b. Remove I/O Enclosure or Components.
 - c. Remove Storage Enclosure or Drives.
 - d. Remove Expansion Rack.
7. Network Setting: Manage or perform changes in the HMC IP address:
 - a. Query/Change IP Address Range.
 - b. Convert SFI To Static IP
 - c. Test Network Connectivity.
 - d. View the Network Topology.
8. Notification setting: Manage notification settings:
 - a. Test Problem Notification.
 - b. Transmit Service Information.
 - c. Manage Serviceable Event Notification.
 - d. View or Change Heartbeat Configuration.
9. HMC1 & HMC2 Management: Manage and change the HMC:
 - a. View Management Console Events.
 - b. Backup Critical Data.
 - c. HMC Rebuild and Recovery.
 - d. Save Upgrade Data.
 - e. Re-harvest HMC Vital Product Data (VPD).
 - f. Query Activate/HMC Role Selection.
 - g. Add Managed System.
 - h. Discover Storage Facility.
 - i. Format Media.
 - j. List Files on Removable Media.

10. Node0 & Node1 Management: Manage and change the Node, controller, or servers:

- a. Set No-rsStart.
- b. Reset No-rsStart.
- c. Launch Advanced System Management (ASM).
- d. AIX Command Processing.
- e. Change/Show LPAR State.
- f. CEC Power Control.
- g. Display CEC Drive Status.
- h. Rebuild CEC Hard Drive.
- i. Advance Utilities:
 - i. Open Terminal Window.
 - ii. Close Terminal Window.
 - iii. Backup Partition Profile.
 - iv. Restore Partition Profile.
 - v. Rebuild Managed System Information.
 - vi. Service Processor Status.
 - vii. Rest or Remove Connections.
 - viii. Reference code History.
 - ix. View License.
 - x. Identify LED.
 - xi. Test LED.

11. Storage System: Manage and change the storage facility:

- a. View Storage Facility State.
- b. Reset Service Intent.
- c. PCIe Graphic Analysis.
- d. Change/Show SFI State.
- e. View/Change Processor and Memory Allocation (Variable Image).
- f. Secure Data Overwrite.
- g. Discontinue Storage Facility.
- h. Advanced Utilities:
 - i. View/Reset Attention Indicators.
 - ii. View Device RM Harvest Phase.
 - iii. View Hardware Topology.
 - iv. View Storage Facility Power Status.
 - v. View Storage Facility Resources States.

12. Service Information Center: Open the service information center.

13. Service GUI (WUI): Open the WUI/HMC console.

9.3.2 Storage Management help functions

The DS GUI provides access to comprehensive help functions. The help functions help you use the GUI and provide in-depth details about the overall DS8000 Storage System and its functions. To access the help contents, click the **Help** icon, and then select **Help Contents** to open a separate window, as illustrated in Figure 9-14.

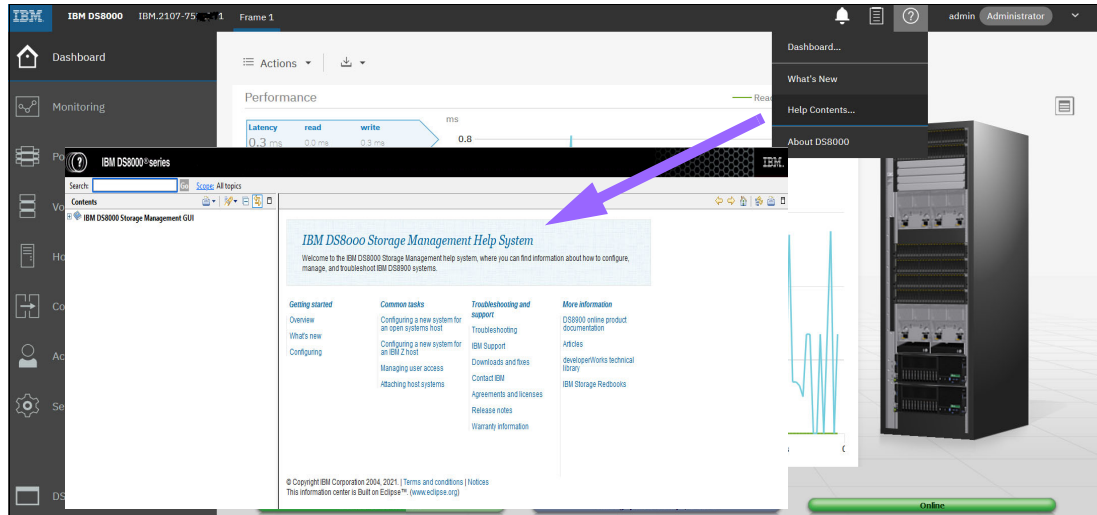


Figure 9-14 Storage Management Help System: IBM Documentation

In IBM Documentation, you can discover introductory information about the DS8900F architecture, features, and advanced functions. You can also learn about the available management interfaces and tools, and troubleshooting and support.

You can obtain more information about using the DS GUI for common tasks:

- ▶ Logically configuring the storage system for open systems and IBM Z attachment
- ▶ Managing user access
- ▶ Attaching host systems

IBM Documentation also provides links to external links for more information about IBM storage systems, and other related online documentation.

For more information, see [IBM DS8900 documentation](#).

9.4 System configuration overview

After the DS8900F administrator configures the initial system, the administrator can configure extra system functions according to the storage requirements.

9.4.1 Network settings

You can change the customer network settings for the DS8900F on the tabs that are described in the following sections.

Ethernet Network

The network settings for both HMCs are performed by IBM Support personnel during system installation. To modify the HMC network information postinstallation, click **Settings** → **Network** → **Ethernet Network**, as shown in Figure 9-15.

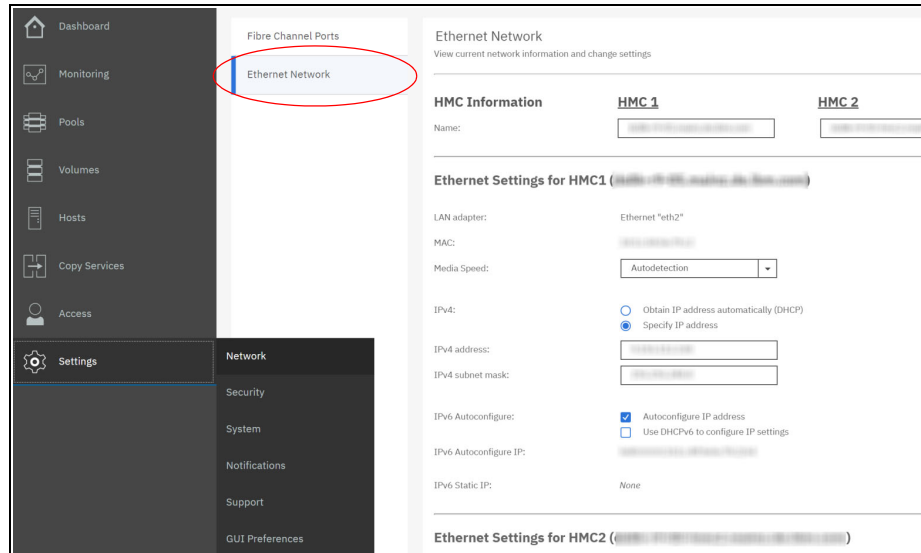


Figure 9-15 Ethernet Network settings

Fibre Channel ports

From the Fibre Channel Ports tab of the Network Settings window, which is shown in Figure 9-16 on page 256, you can configure the protocol that is used by the FC ports or view the properties of those ports:

- ▶ **Modify Protocol**
Set the protocol that is used for connecting to a host or another storage system.
- ▶ **Modify Endpoint Security**
Set a new IBM Fibre Channel Endpoint Security level: Disable, Enable, or Enforced.
- ▶ **Refresh Security**
Port refresh causes the port to be taken offline for a few seconds and put back online.
- ▶ **Logged in WWPNS**
Select this option to view details about connections that are logged in to an FC port.
- ▶ **Error Rates**
View the error rates of an FC port. For more information, see 9.15, “Fibre Channel error rate statistics” on page 336.
- ▶ **Performance**
Use this option to view graphs showing key performance metrics for FC ports, such as IOPS, Latency, and Bandwidth. For more information, see “Creating FC port performance graphs” on page 332.

- ▶ Properties

View the properties of an FC port, such as the state and protocol.

- ▶ Export Fibre Channel Ports information

Click the **Download** icon next to the **Actions** menu to create and download a CSV file that contains all the information for FC ports that are shown in the file, including their worldwide port names (WWPNs).

Figure 9-16 shows the FC ports window with all available options that are listed.

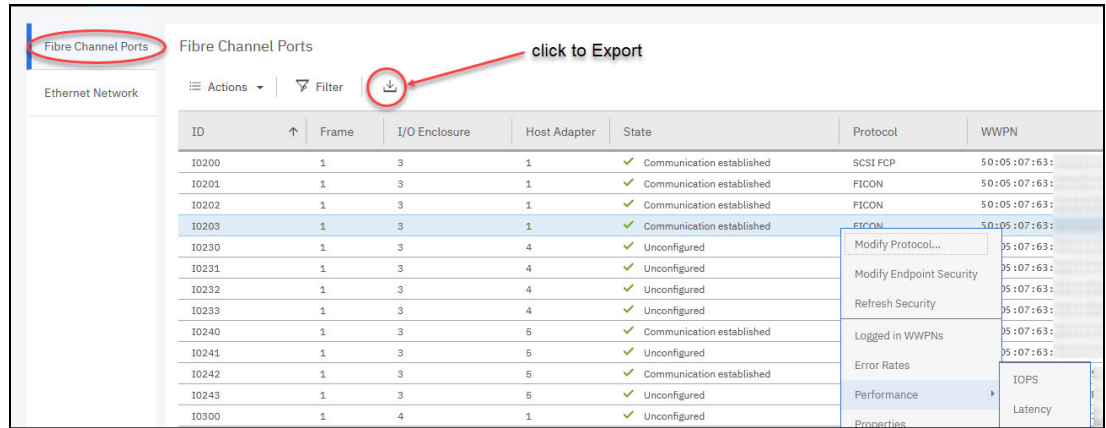


Figure 9-16 Fibre Channel Ports settings

Note: Exporting the FC port information does not produce the comprehensive report that is available in the FC connectivity report.

9.4.2 Security settings

To protect and safeguard data on the system, the DS8900F offers the following encryption features:

- ▶ Data at rest encryption: Disk-based encryption on the flash drives is combined with an enterprise-scale key management infrastructure to encrypt the data at rest in the drives.
- ▶ IBM Fibre Channel Endpoint Security: Protects data in flight between an IBM Z host and the DS8900F storage system by controlling access and encrypting data that is transferred over a storage area network (SAN).
- ▶ Local password rules to accommodate any specific company imposed rules for password assignments.
- ▶ Communications certificate on the HMC to enable HTTPS connections with the storage system.

Use these settings to configure the security settings for your DS8900F system.

Data-at-rest encryption

To enable data-at-rest encryption, select the **Settings** icon from the DS GUI navigation menu on the left. Click **Security** to open the Security window, and click the **Data at Rest Encryption** tab, as shown in Figure 9-17 on page 257.

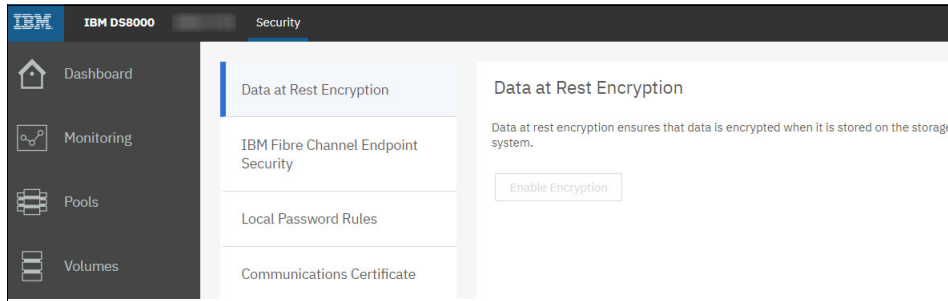


Figure 9-17 Enabling data-at-rest encryption

You can define a custom certificate for communication between the encryption key servers (typically IBM Security Guardium Key Lifecycle Manager) and the storage system.

A system-generated Gen-2 (National Institute of Standards and Technology (NIST) SP 800-131a compliant) or Gen-3 certificate can be updated to a customer-generated certificate. The custom certificate must meet all the requirements or the update fails.

Important: If you plan to activate data-at-rest encryption for the storage system, ensure that the encryption license is activated and the encryption group is configured before you begin any logical configuration on the system. After the pools are created, you cannot disable or enable encryption.

If the DS8900F was ordered with the Local Key Management feature, then the DS8900F manages the key group. Local Key Management can be set up only by using the DS CLI. For more information, see *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security (DS8000 Release 9.2)*, REDP-4500.

IBM Fibre Channel Endpoint Security

IBM Fibre Channel Endpoint Security protects data in flight between an IBM Z host and the DS8900F storage system by controlling access and encrypting data that is transferred over a SAN. For more information about IBM Fibre Channel Endpoint Security, see *IBM Fibre Channel Endpoint Security for IBM DS8900F and IBM Z*, SG24-8455.

To enable IBM Fibre Channel Endpoint Security, select the **Settings** icon from the **DS GUI navigation** menu on the left. Click **Security** to open the Security window, and click the **Fibre Channel Endpoint Security** tab, as shown in Figure 9-18.

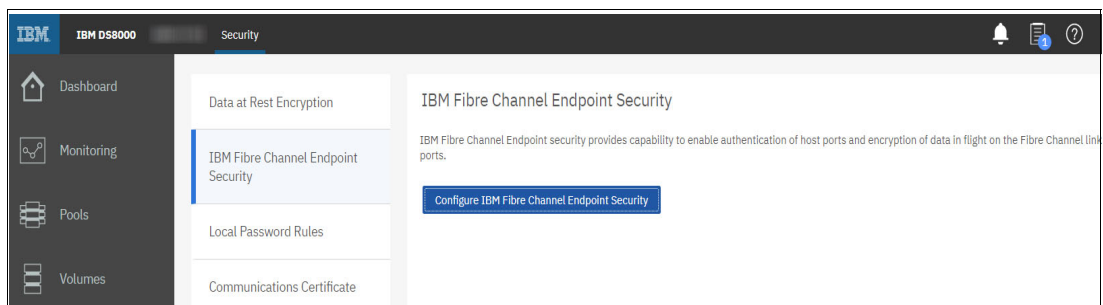


Figure 9-18 Configuring IBM Fibre Channel Endpoint Security

Local password rules

Many companies have their own security policies for passwords, and want to implement the same policies for their DS8900F storage system.

To implement password rules, complete these steps (see Figure 9-19):

1. From the system window, click the **Settings** icon.
2. Click **Security** to open the Security window.
3. Click the **Local Password Rules** tab.

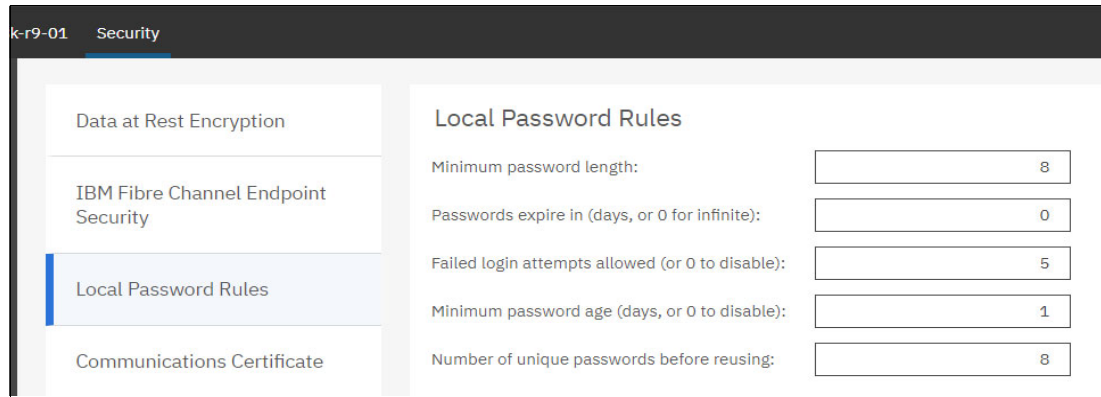


Figure 9-19 Local Password Rules

Communications Certificate

The **Communications Certificate** tab of the Security Settings window can be used to assign or create an encryption certificate for each HMC with HTTPS connections to the storage system. You can also create certificate signing requests (CSRs), import existing certificates, create self-signed certificates, and view the certificate information for each HMC, as shown in Figure 9-20.

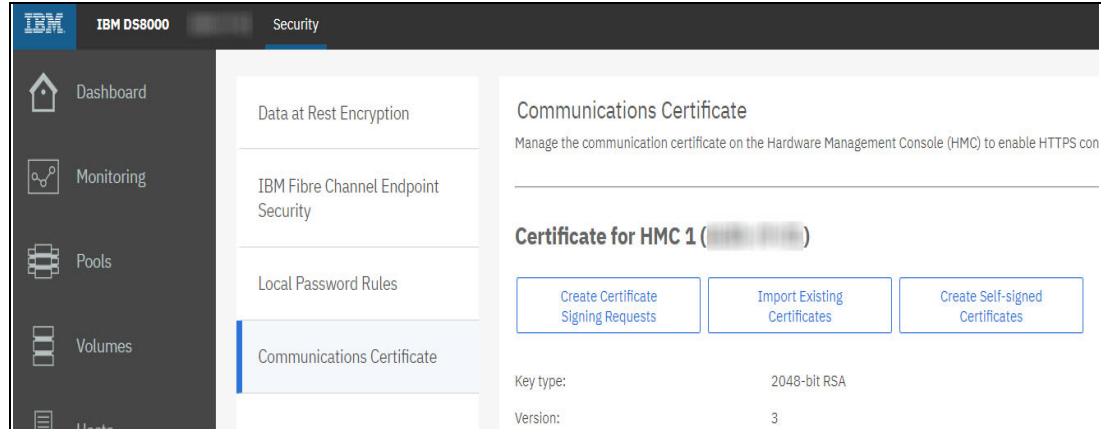


Figure 9-20 Communication Certificates

The **Create Certificate Signing Request** button is used to generate a CSR that is sent to a certificate authority (CA) for verification. As shown in Figure 9-21 on page 259, the necessary information to include in the CSR are the HMC fully qualified domain name (FQDN), organization details, the length of time that the certificate must be valid, and an email address.

Create Certificate Signing Request

 HMC 1 (XXXXXXXXXX)

Key type: 2048-bit RSA

HMC DNS hostname:

Organization:

Organization unit:

Country:
 ▼

State or province:
 ▼

City or locality:

Number of days until expiration:

Email address:

Figure 9-21 Certificate signing request

After the CSR file is created, you can download that file for processing with your trusted CA.

The extra two options that are available here for secure communications to the DS8900F system are to import an already provided CA certificate from your security group within your organization, or to create a self-signed certificate.

9.4.3 System Settings

The next option under Settings is System, which includes system settings for Licensed Functions, Easy Tier, zHyperLink, Date and Time, and Advanced.

Licensed Functions

You can display all the installed licensed functions and activate new function keys from this menu, as shown in Figure 9-22.

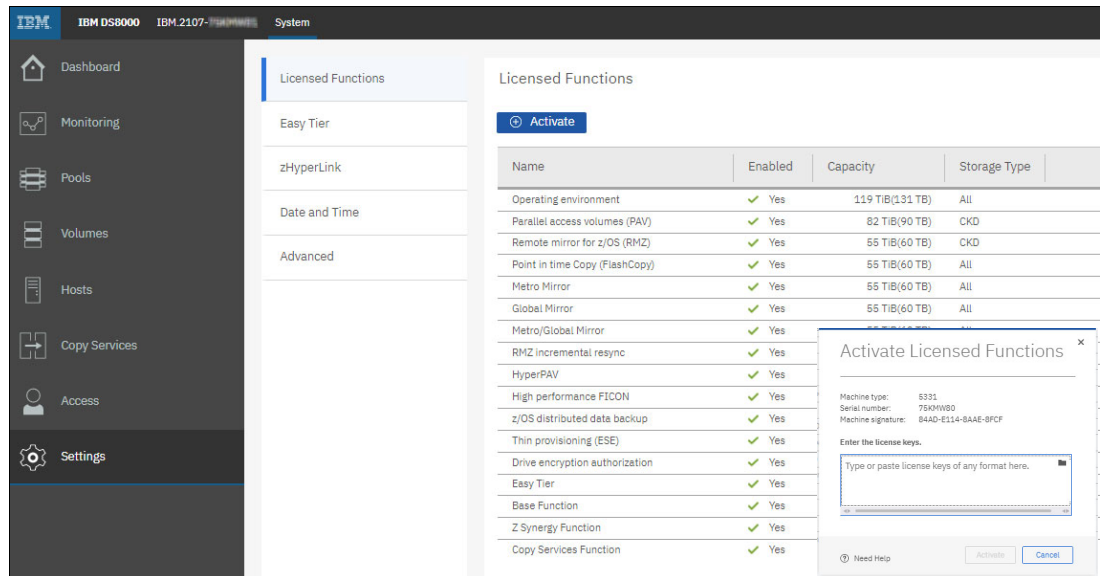


Figure 9-22 Licensed Functions settings

Easy Tier settings

To configure Easy Tier controls and advanced settings on your DS8900F system, select the **Settings** icon from the navigation menu on the left. Click **System** to open the System window, and click the **Easy Tier** tab to open the Easy Tier Settings window, as shown in Figure 9-23.

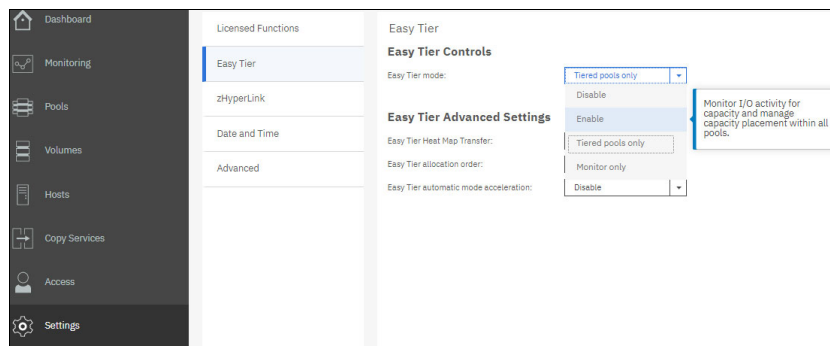


Figure 9-23 Easy Tier Settings

The following settings are available:

- ▶ **Easy Tier mode:** The available options are **Enable**, **Tiered pools only**, **Disable**, **Monitor only**, or **Disable**. When this setting is configured to **enable**, Easy Tier monitors I/O activity for capacity in all pools, and manages capacity placement within them.
- ▶ **Easy Tier Heat Map Transfer (HMT):** Use this setting to maintain application-level performance at the secondary site of a DS8000 by transferring the Easy Tier information to the secondary site.
- ▶ **Easy Tier Allocation order:** Specify the allocation order that is used by Easy Tier to select the drive classes when allocating capacity in a pool.
- ▶ **Easy Tier Automatic mode acceleration:** Use this setting to temporarily accelerate data migration by Easy Tier.

For more information about Easy Tier settings, see *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667.

zHyperLink

zHyperLink is a short-distance link technology that complements Fibre Channel connection (IBM FICON) technology to accelerate I/O requests that are typically used for transaction processing. It consists of point-to-point connections for random reads and writes, and provides up to 10 times lower latency than High-Performance FICON for IBM Z (zHPF). You can set it to **Enabled**, **I/O Read Enabled**, **I/O Write Enabled**, or **Disabled**, as shown in Figure 9-24.

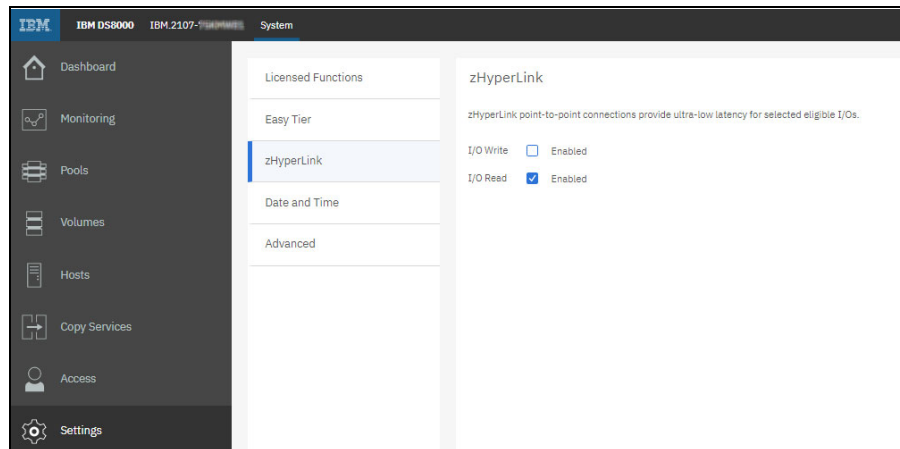


Figure 9-24 zHyperLink settings

Note: To take advantage of zHyperLink in DS8000, ensure that CUIR support (under IBM Z) is enabled.

For more information, see *Getting Started with IBM zHyperLink for z/OS*, REDP-5493.

Date and Time

You can enter the date and time manually, or specify an external NTP server to provide the date and time, as shown in Figure 9-25.

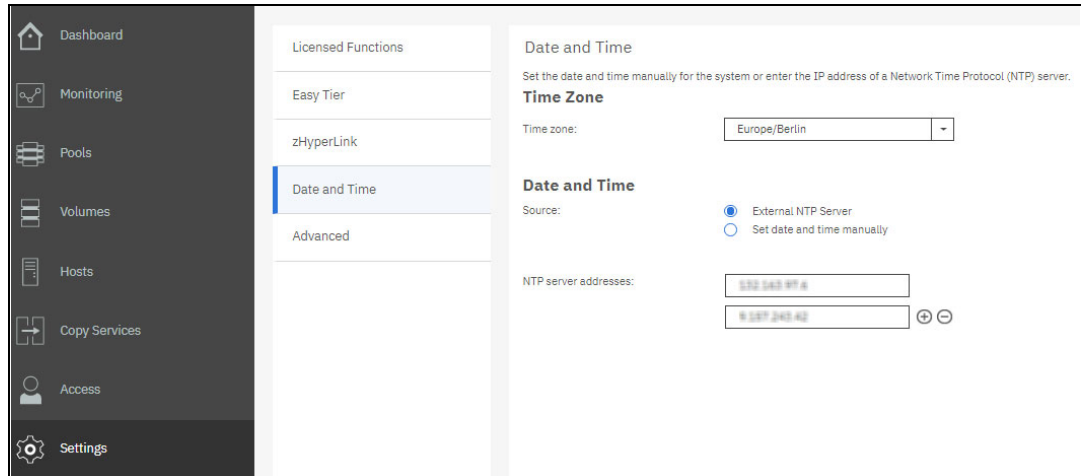


Figure 9-25 Date and Time settings

Advanced settings

From the Advanced tab of the System Settings window, you can configure system-wide options, such as power management, CS, IBM Z features, and the DS Open application programming interface (API).

The Advanced settings window is shown in Figure 9-26.

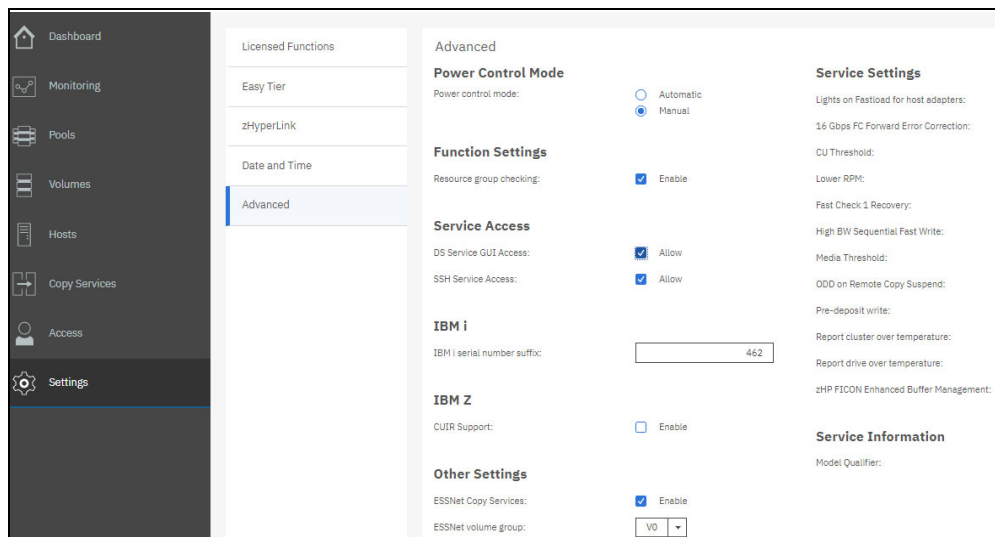


Figure 9-26 Advanced settings

► Power control mode

You can determine how to control the power supply to the storage system. From the System window, click the **Settings** icon. Click **System** to open the System window. Click the **Advanced** tab to open the window to manage Power control mode (as shown in Figure 9-26 on page 262). The following options are available:

- Automatic: Control the power supply to the storage system through the external wall switch.
- Manual: Control the power supply to the storage system by using the **Power Off** action on the System window.

► Function settings

The Resource Group Control option is available in the Function Settings section. It allows a storage administrator to specify which users can perform certain logical configuration actions, such as create or delete volumes in a pool.

► Service Access

The following options are available in the Service Access section:

- DS Service GUI Access.
Allows authorized IBM SSRs to access the DS Service GUI.
- SSH Service Access.
Allows authorized IBM SSRs to access the Secure Shell (SSH) CLI on the HMC.

► IBM i

The following option is available in the IBM i section:

IBM i serial number suffix: Enter the IBM i serial number suffix to avoid duplicate logical unit number (LUN) IDs for an IBM i (AS/400) host. Restart the storage system to assign the new serial number.

► IBM Z

The following option is available in the IBM Z section:

CUIR Support: Enables control unit initiated reconfiguration. This option allows automation of channel path quiesce and resume actions during certain service actions. It eliminates the requirement for manual actions from the host.

► Other settings

The following options are available in the Other Settings section:

- ESSNet CS.
Enables the ESSNet user interface to manage CS on the storage system.
- ESSNet volume group.
Selects the ESSNet user interface to manage the volume group with CS.
- Host precheck.
Enables FB and CKD volume delete protection.
- Device Threshold.
Sets the threshold level for IBM Z at which the system presents a service information message (SIM) to the operator console for device-related errors. Device threshold levels are the same type and severity as control unit threshold settings:
 - 0: Service, Moderate, Serious, and Acute (all)
 - 1: Moderate, Serious, and Acute

- 2: Serious and Acute
 - 3: Acute
- Full Page Protection.
Enables the ability to ensure that the atomicity of a database page-write is maintained.
 - PPRC Path Recovery.
Enables the storage system to monitor PPRC paths for signs of failure. If a path fails, it is placed into a degraded state in which it is used minimally until the problem stops. If this setting is disabled, paths are never put into a degraded state.
 - Present SIM data to all hosts.
Enables SIMs to be sent to all or to only the first attached IBM Z logical partition (LPAR), and makes an I/O request to the logical system or logical volume. This setting applies to IBM Z environments only.
 - Enhanced Information Unit Pacing.
Enables increased write performance of large writes at long distances and improvement of z/OS Global Mirror (zGM) initial copy performance because the channel can send more read track commands to the primary storage system.
- ▶ Automatic code management
Automatically download and preinstall the recommended code level to update the code level. The default setting is **Enable**.
 - ▶ Service settings
You can view the service settings but they cannot be changed.

9.4.4 Notifications settings

You can configure and manage Call Home and Syslog settings for the system from the DS GUI by selecting **Notifications** from the **Settings** menu, as shown in Figure 9-27.

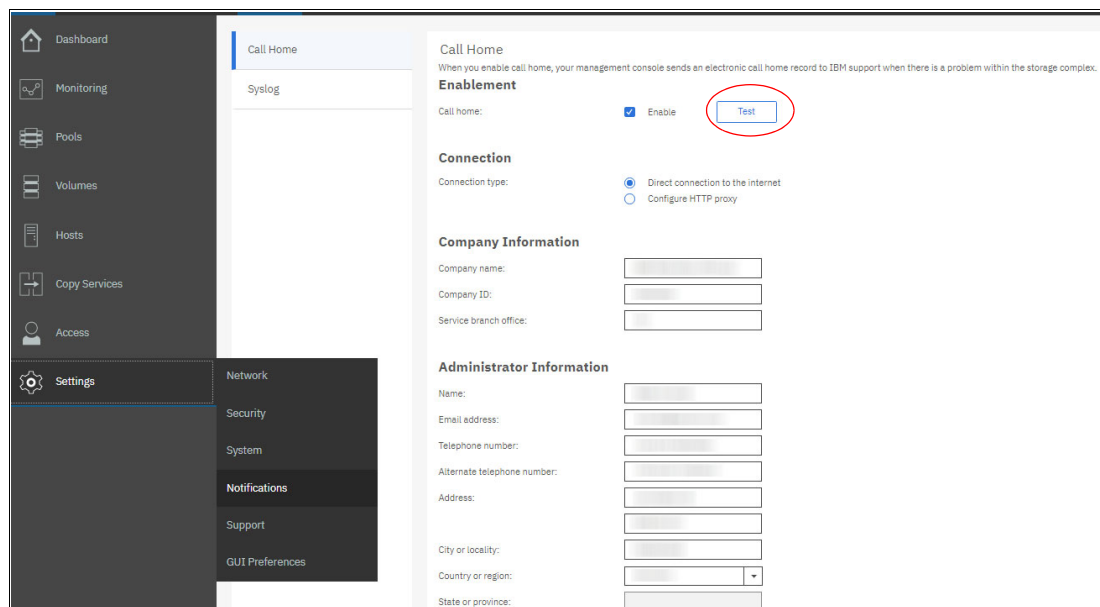


Figure 9-27 Notifications Settings window

Call Home

The DS8900F uses the Call Home feature to report serviceable events to IBM. To ensure timely action from IBM Support personnel for these events, it is important to enable and properly configure Call Home on the system.

When enabling Call Home for the first time, you must accept the Agreement for Service Program when presented. Enter your **Company Information**, **Administrator Information**, and **System Information** details. Finally, after completing the setup, you can test the Call Home feature by clicking **Test**, as shown in Figure 9-27 on page 264.

Syslog

The Syslog window displays the syslog servers that are configured to receive logs from the DS8900F system. A user with an administrator role can define, modify, or remove up to eight syslog target servers. Each syslog server must use the same TLS certificate. Events such as user login and logout, commands that are issued by an authorized user by using the DS GUI or DS CLI, and remote access events are forwarded to syslog servers. Additionally, events in the RAS audit log and Product Field Engineer (PFE) actions are also forwarded to the syslog servers. Messages from the DS8900F are sent by using facility code 10 and severity level 6.

To configure one or more syslog servers, complete these steps:

1. Click **Settings** → **Notifications**.
2. On the Notifications window, select **Add Syslog Server**. You receive a warning to enable TLS first before adding any syslog server, as shown in Figure 9-28.




Figure 9-28 Warning to enable TLS

Note: A DS8900F server must use TLS for its communications with the syslog server. To configure TLS, the customer must generate their own trusted certificate for the DS8900F syslog process with the CA and import the trusted CA file, signed machine (in this case the HMC and syslog process) syslog server certificate file, and key file, as shown in Figure 9-29.

For more information about the setup of the SYSLOG server with TLS, see [Encrypting Syslog Traffic with TLS \(Secure Sockets Layer\) \(SSL\)](#).

The process involves external entities such as your trusted CA and potentially the use of the **openssl** command to retrieve the syslog server generated key if it is not already provided by the CA.

The files that are entered into the fields that are shown in Figure 9-29 are:

- ▶ CA Certificate (ca.pem)
- ▶ HMC Signed Certificate (cert.pem)
- ▶ HMC Key (key.pem)

3. To enable TLS, in the Certificates area, click **Enable TLS**, as shown in Figure 9-29.

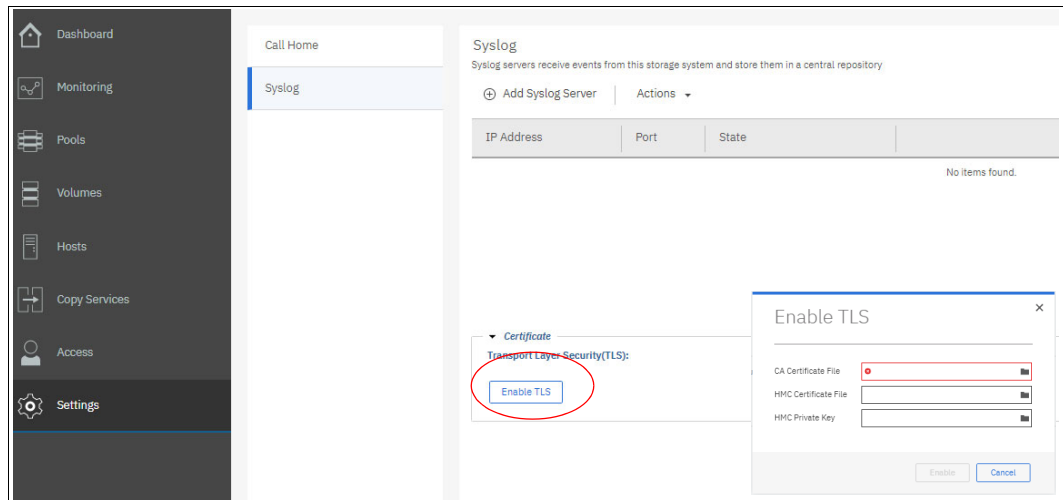


Figure 9-29 Enable TLS window

4. In the Enable TLS window, browse for the following certificate files on your local machine:
 - The CA certificate file (Example: ca .pem).
 - The syslog communications certificate file, which is signed by the CA. (Example: hmc .pem).
 - The extracted Private Key file, which is the private key for the storage system. (Example: key .pem).
5. Click **Enable** to complete the TLS configuration.
6. To add a syslog server, click **Add Syslog Server**, as shown in Figure 9-30, and provide the following parameters:
 - **IP Address:** The IP address of the external syslog server.
 - **Port:** The TCP port for the external syslog server (the default is 514).
7. After you review the details, click **Add** to create the syslog server entry.

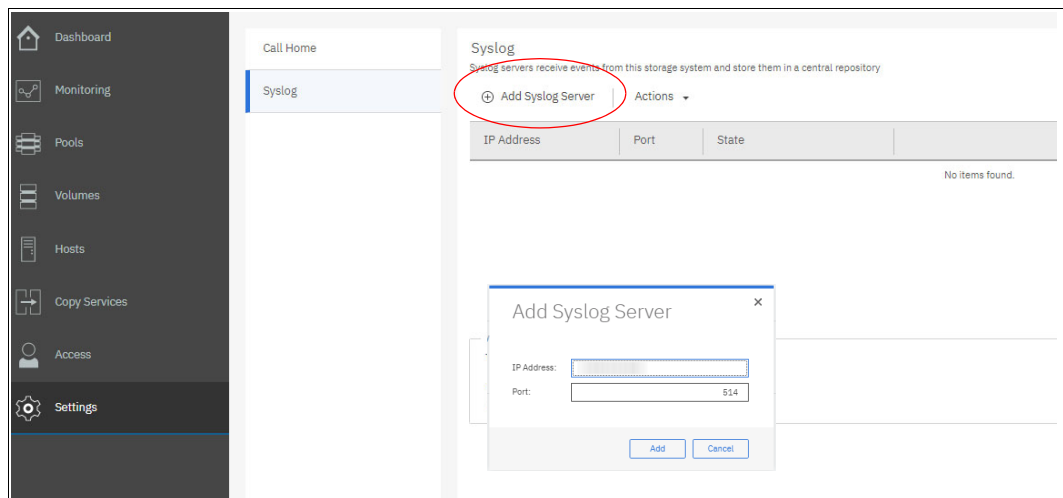


Figure 9-30 Syslog settings window

- After the required syslog servers are created, you can **Modify**, **Test**, **Activate**, **Deactivate**, and **Remove** a selected syslog server, as shown in Figure 9-31.

Note: To enable TLS, all existing syslog servers must be deleted first. Then, you can enable TLS and create the syslog servers.

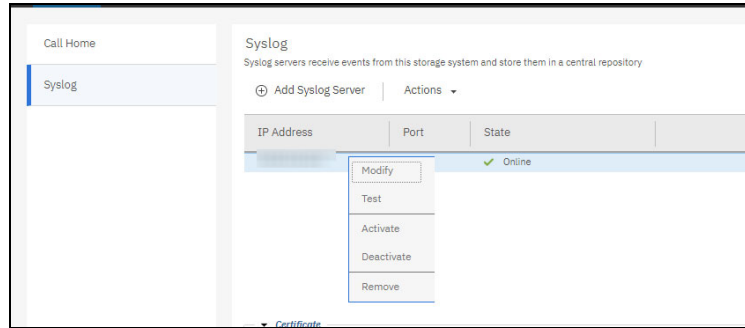


Figure 9-31 Modify Syslog configuration

9.4.5 Support settings

Use this section to configure various support settings.

IBM Remote Support Center

On the Support settings window, you can configure service access to the HMC to allow RSC to access the HMC for problem determination, as shown in Figure 9-32.

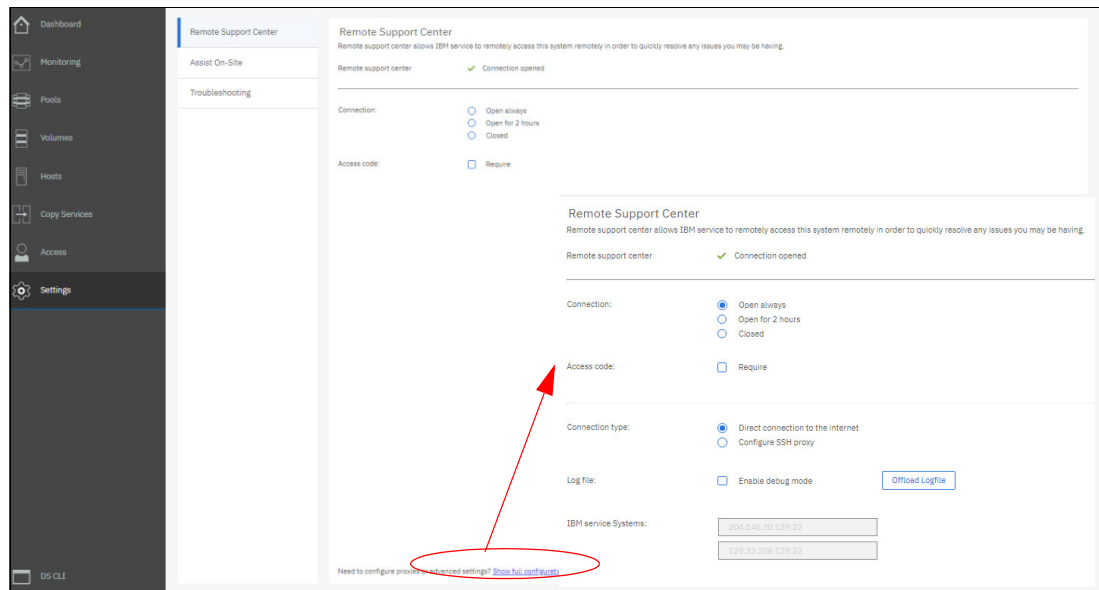


Figure 9-32 IBM Remote Support Center settings

You can configure the RSC access to stay open continuously, close 2 hours after RSC logs off, or keep it closed. You can require IBM service to use an access code for remote support connections with the HMC on your storage system. Click **Generate** to generate an access code or enter your own access code. The access code is case-sensitive and must be fewer than 16 characters.

To configure the RSC, click **Show Full Configuration**. You can select the connection type, define your proxy information, and **Offload Logfile**. **Offload Logfile** downloads the RSC log file, which lists events and actions for the storage system. Select **Enable debug mode** to generate detailed RSC log files. **IBM Service Systems** lists the IP addresses of the RSC servers that you must configure your firewall to allow access to these addresses.

Assist On-site

If AOS is used for an IBM Support connection to the HMC, you can **Start**, **Stop**, or **Restart** the AOS service from the GUI, as shown in Figure 9-33.

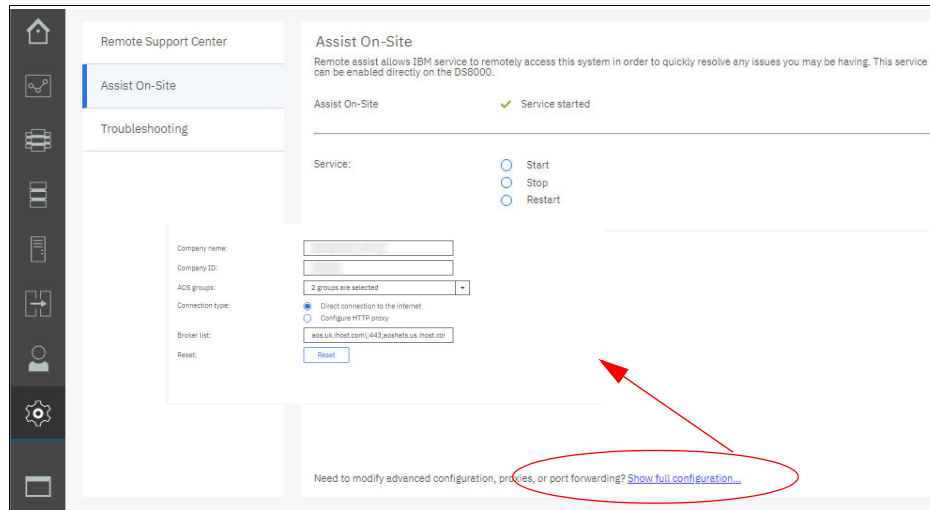


Figure 9-33 AOS settings and full configuration

To configure AOS, click **Show Full Configuration** and enter the required settings, as shown in Figure 9-33.

Troubleshooting

Use the **Troubleshooting** tab to perform actions that resolve common issues with your storage system:

- ▶ Restart HMCs

If there are connectivity issues with the storage management software (DS GUI, DS CLI, IBM Copy Services Manager, or IBM Spectrum Control), click **Restart HMC**. You can also use this feature to restart an HMC after you modify the settings of the HMC.

- ▶ Refresh GUI Cache

If there are inconsistencies between what is displayed in the DS GUI and the DS CLI or IBM Spectrum Control, click **Refresh GUI Cache**.

- ▶ Reset Communications Path

To restart the web servers and communication paths that are used by IBM ESSNI, click **Reset Communications Path**.

Figure 9-34 shows the Troubleshooting tab.

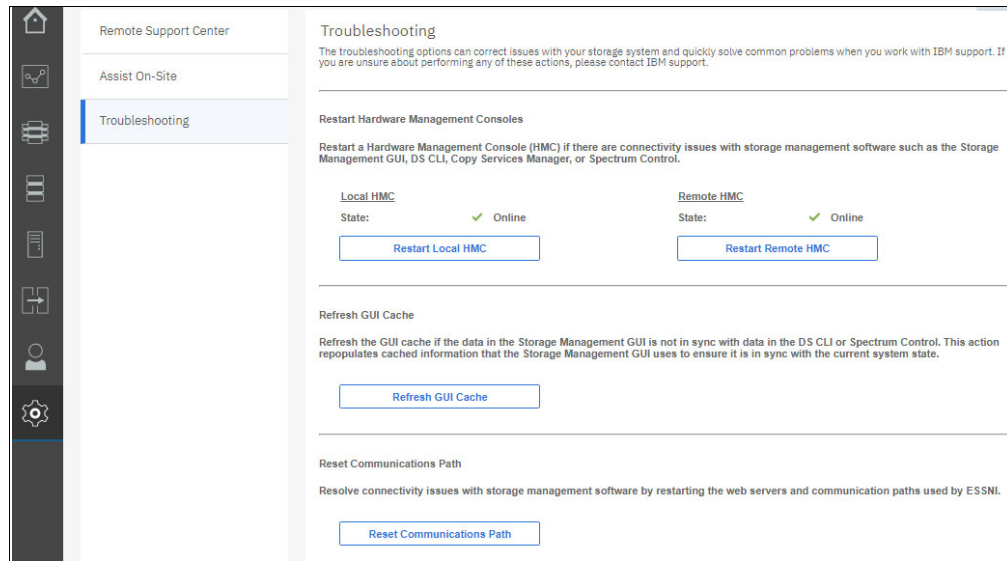


Figure 9-34 Troubleshooting tab

GUI Preferences

Use the GUI Preferences tab that is shown in Figure 9-35 to set the following options for the DS GUI:

- ▶ Login Message

With an administrator role, you can enter a message that is displayed when users log in to either the DS GUI or the DS CLI.

- ▶ General GUI settings

On the General tab of the GUI Preferences window, you can set the default logout time for the DS GUI.

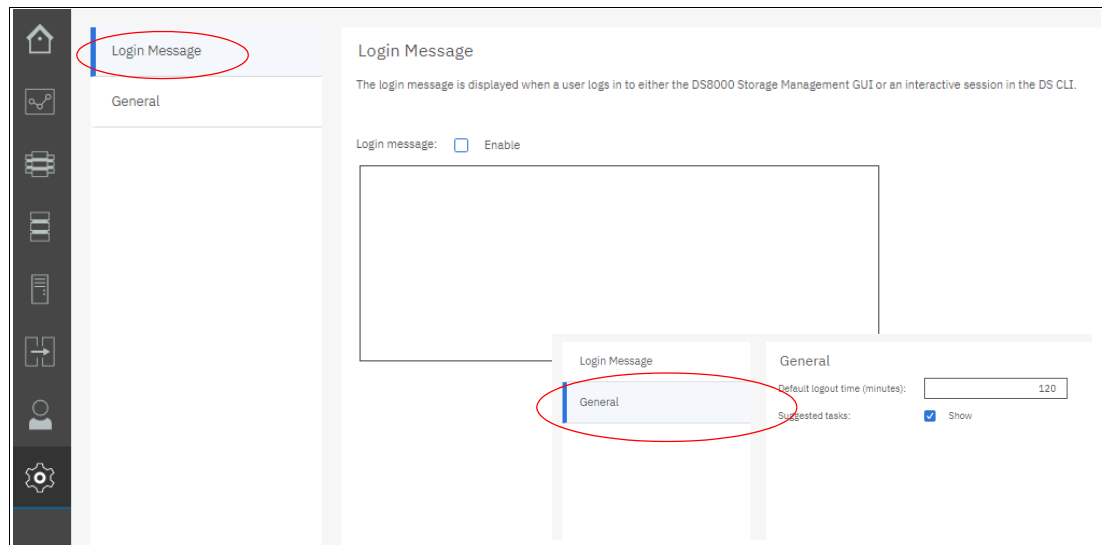


Figure 9-35 Login message

9.5 Logical configuration overview

The logical configuration of the DS8900F storage system begins with managed arrays and the creation of storage pools.

When the storage pools are created, arrays are first assigned to the pools, and then volumes are created in the pools. FB volumes are connected through host ports to an open system host. CKD volumes require LSSs to be created 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 performing 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. Multiple pools can be created to isolate workloads.

When you create a pool pair, all available arrays can be assigned to the pools, or the choice can be made to manually assign them later. 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 the storage connects to an IBM Z host, you must create the LSSs before you create the CKD volumes.

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

When you plan your configuration with the DS8900F, all volumes, including standard provisioned volumes, use metadata capacity when they are created, which causes the usable capacity to be reduced. The 1 (gibibyte) GiB extents that are allocated for metadata are subdivided into 16 mebibyte (MiB) subextents. The metadata capacity of each volume that is created affects the configuration planning.

If the volumes must connect to an IBM Z host, the next steps of the configuration process are completed on the host. For more information about logically configuring storage for IBM Z, see 9.7, “Logical configuration for Count Key Data volumes” on page 292.

If the volumes connect to an open system host, map the volumes to the host, and then add host ports to the host and map them to FC ports on the storage system.

FB volumes can accept I/O only from the host ports of hosts that are mapped to the volumes. Host ports are zoned to communicate only with certain FC ports on the storage system. Zoning is configured either within the storage system by using FC port masking, or on the SAN. Zoning ensures that the workload is spread correctly over FC ports and that certain workloads are isolated from one another.

Host configuration is simplified by the DS8900F microcode. Host ports are now automatically updated and host mappings can be performed during the volume creation step of the logical configuration. In addition, host port topology can be safely changed by using the DS GUI and DS CLI. New host commands are available for DS CLI to make, change, delete, list, and show a host connection. For more information, see Chapter 10, “IBM DS8900F Storage Management Command-line Interface” on page 339.

Note: Deleting a pool with volumes is available in the GUI. A warning is displayed, and the user must enter a code that is presented by the DS8900F to confirm the delete. A “force deletion” option is also available. For more information, see Figure 9-88 on page 306.

9.6 Logical configuration for open systems volumes

This section describes the logical configuration for Fixed-Block (FB) volumes for open systems hosts. It covers the following topics:

- ▶ Simple open systems (FB) logical configuration flow
- ▶ FB pool creation
- ▶ Quick FB volume creation
- ▶ Advanced FB volume creation
- ▶ Creation and connection of FB volumes to the open systems hosts

9.6.1 Configuration flow

The following steps provide an overview of the steps that are needed for logical configuration of FB volumes:

1. Create an FB pool pair for open systems hosts.
2. Create the FB volumes.
3. Map to the open system hosts.

9.6.2 Creating FB pools for open systems hosts

For best performance and a balanced workload, create two pools. The DS GUI helps the system administrator to create a balanced configuration by creating pools as a pair. The pools are configured so that one pool of the pair is managed by system node 0 and the other pool of the pair is managed by node 1.

Note: If the requirement is to create a single pool, see “Creating a single pool” on page 277.

To create an FB pool pair, complete these steps:

1. Click the **Pools** icon and select the **Arrays by Pool** option to open the Array by Pool window, as shown in Figure 9-36.

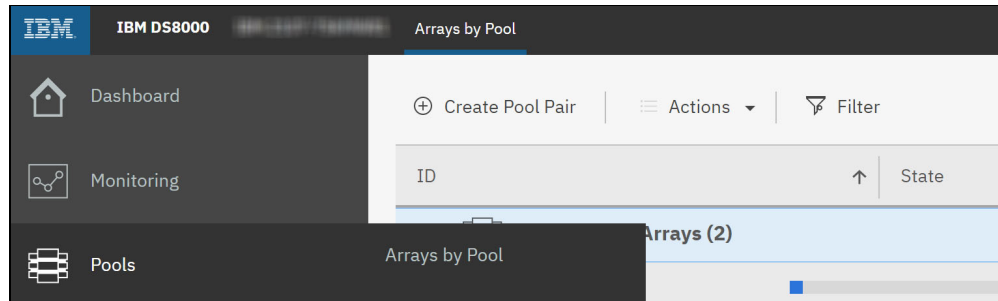


Figure 9-36 Arrays by Pool

2. Click the **Create Pool Pair** tab, as shown in Figure 9-37. The Create Pool Pair window opens.

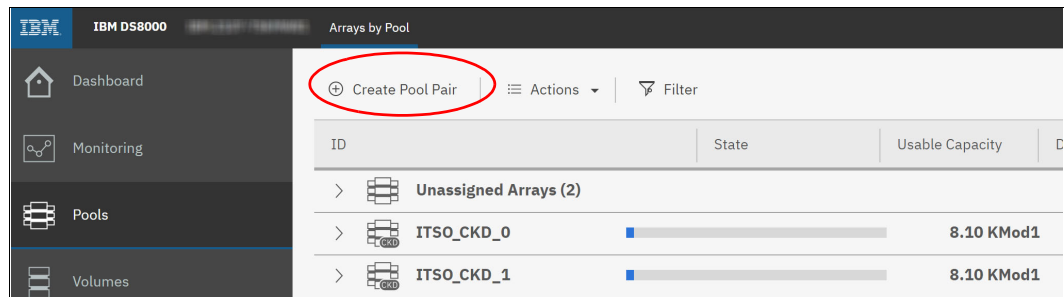


Figure 9-37 Create Pool Pair window

Note: You can automatically assign arrays when creating a pool pair. The arrays are created with the default redundant array of independent disks (RAID) type, RAID 6. To configure other supported raid types, select the **Custom** option under the Create Pool Pair dialog, or assign arrays manually to an existing storage pool from the Unassigned Arrays. RAID 5 needs a Request for Price Quotation (RPQ), but it is not recommended. For more information, see “Creating Fixed-Block pools: Custom” on page 274.

3. Specify the pool pair parameters, as shown in Figure 9-38 on page 273:
 - Storage type: Ensure that **Open Systems (FB)** is selected.
 - Name prefix: Add the pool pair name prefix. A suffix ID sequence number is added during the creation process.




Figure 9-38 Creating an FB pool pair and assigning arrays

4. Select from the listed drive types and select the number of arrays for each drive type that you want to assign to the pool pair.

Important: The number of specified arrays must be even. Trying to specify an odd number results in a message that states “Arrays must be spread evenly across the pool pair”. The GUI increases the number of arrays by one to achieve an even number.

5. When pool pair parameters are correctly specified, click **Create** to proceed. Figure 9-39 shows a pool pair that is created and assigned arrays.

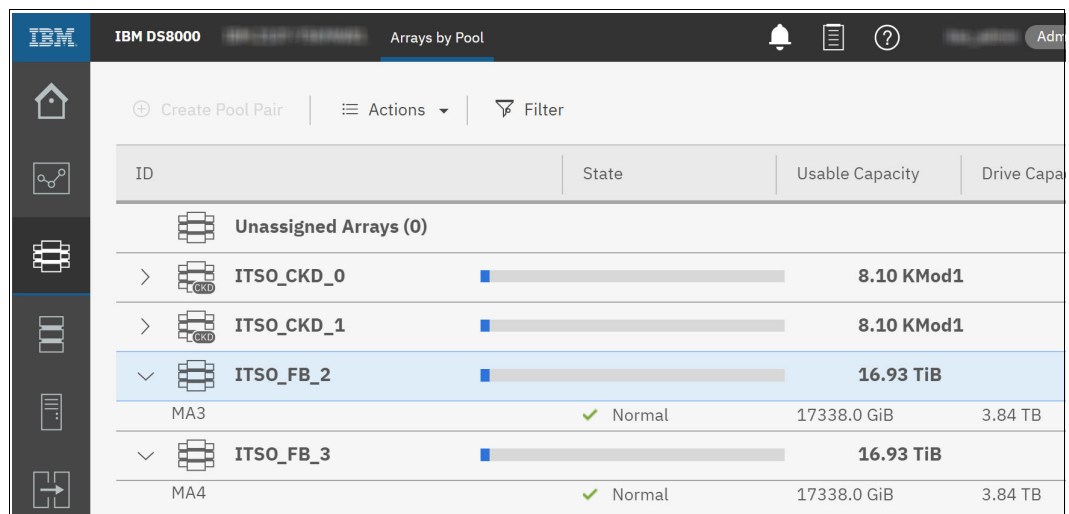


Figure 9-39 Pool pair that is created with assigned arrays

Creating Fixed-Block pools: Custom

To specify the extent size or RAID level for the arrays at pool creation time, select the **Custom** option from the Create Pool Pair dialog.

Available options for extent size are 1 GiB (large), or 16 mebibytes (MiB) (small). **Small extent size** is the preferred option because it provides better capacity utilization. For large systems that use Easy Tier, it might be preferable to use large extents. For an in-depth description about large and small extents, see Chapter 4, “Virtualization concepts” on page 107.

To create a custom FB pool pair, complete these steps:

1. Click the **Pools** icon and select the **Arrays by Pool** option to open the Array by Pool window, as shown in Figure 9-36 on page 272.
2. Click the **Create Pool Pair** tab. The Create Pool Pair window opens.
3. Select the **Custom** option.
4. Specify the pool pair parameters, as shown in Figure 9-40 on page 275:
 - Storage type: Ensure that **Fixed block (FB)** is selected.
 - Name prefix: Add the pool pair name prefix. A suffix ID sequence number is added during the creation process.
 - Extent size: Select **1 GiB** for large extents, or **16 MiB** for small extents.
5. Select from the listed drive types and select the number of arrays for each drive type that you want to assign to the pool pair.
6. Choose the RAID level for the selected arrays. RAID 6 is the recommended and default RAID type for all drives over 1 TB.

Note: RAID 5 is supported only for drives less than 1 TB and requires an RPQ. If selected, you must acknowledge your understanding of the risks that are associated with RAID 5 before continuing. For more information about the supported drive types and available RAID levels, see Chapter 2, “IBM DS8900F hardware components and architecture” on page 25.

7. When the pool pair parameters are correctly specified, click **Create** to proceed, as shown in Figure 9-40 on page 275.

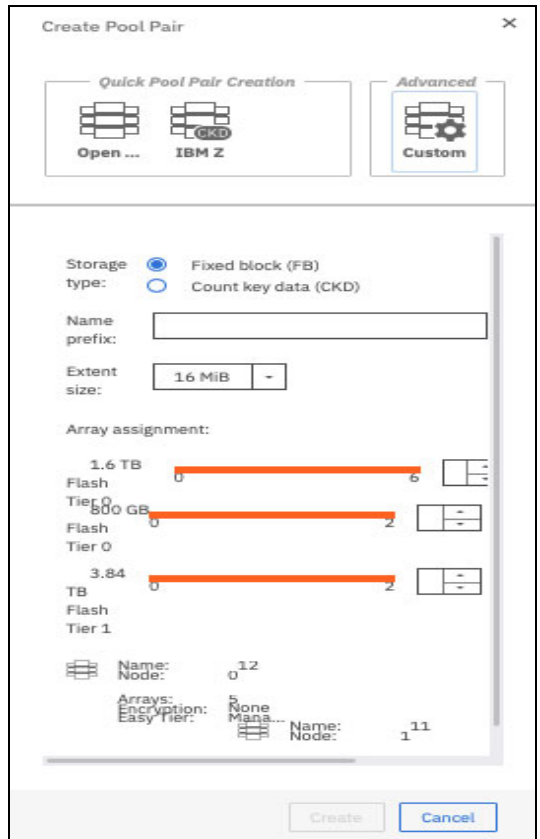


Figure 9-40 Create FB pools (Custom)

Manually assigning arrays to existing pools

The storage administrator can manually unassign or reassign arrays from or to existing pools when the current configuration must be modified, such as when the administrator adds storage capacity. To manually assign arrays, complete these steps:

1. Select an array and click **Assign** or **Reassign**, as shown in Figure 9-41. This action opens the Assign Array window.

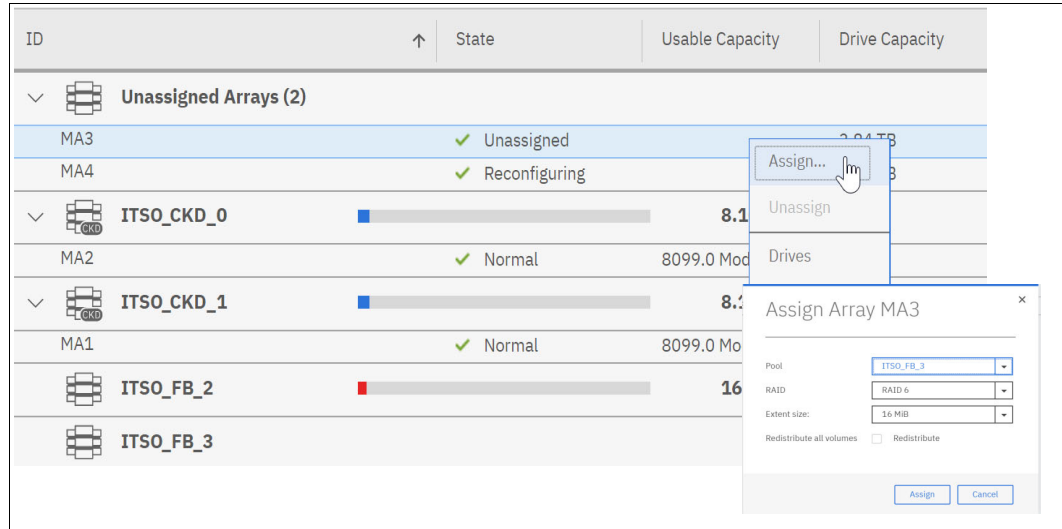


Figure 9-41 Manually assigning an array to an existing pool

2. Select the target pool from the drop-down list, and the RAID level that you want.
3. Select the **Redistribute** checkbox to redistribute all existing volumes across the pool, including the new array.
4. Click **Assign**.

Note: In a pool that is managed by Easy Tier, redistributing volumes across the pool is automatic. This redistribution is called *Dynamic Pool Reconfiguration*. For more information, see *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667.

Creating a single pool

Occasionally, you are required to create a single pool, as opposed to creating a pool pair for balancing a workload. To create a single storage pool, complete these steps:

1. Create a pool pair, as shown in Figure 9-42. However, do not assign any arrays to the new pool pair.




Figure 9-42 Creating an empty pool pair with no assigned arrays

2. Click **Create**. Two pools are created as usual.
3. Choose one of the pools from the recently created pool pair to delete, as shown in Figure 9-43.

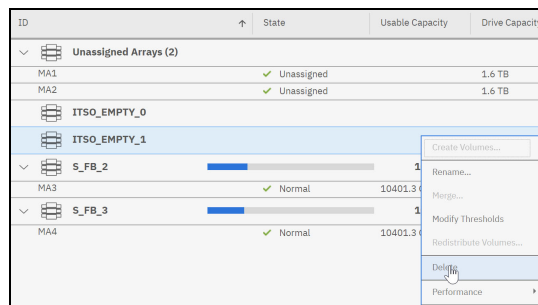


Figure 9-43 Deleting one pool of the pool pair

4. Assign one or more arrays to the single pool, as shown in Figure 9-44.

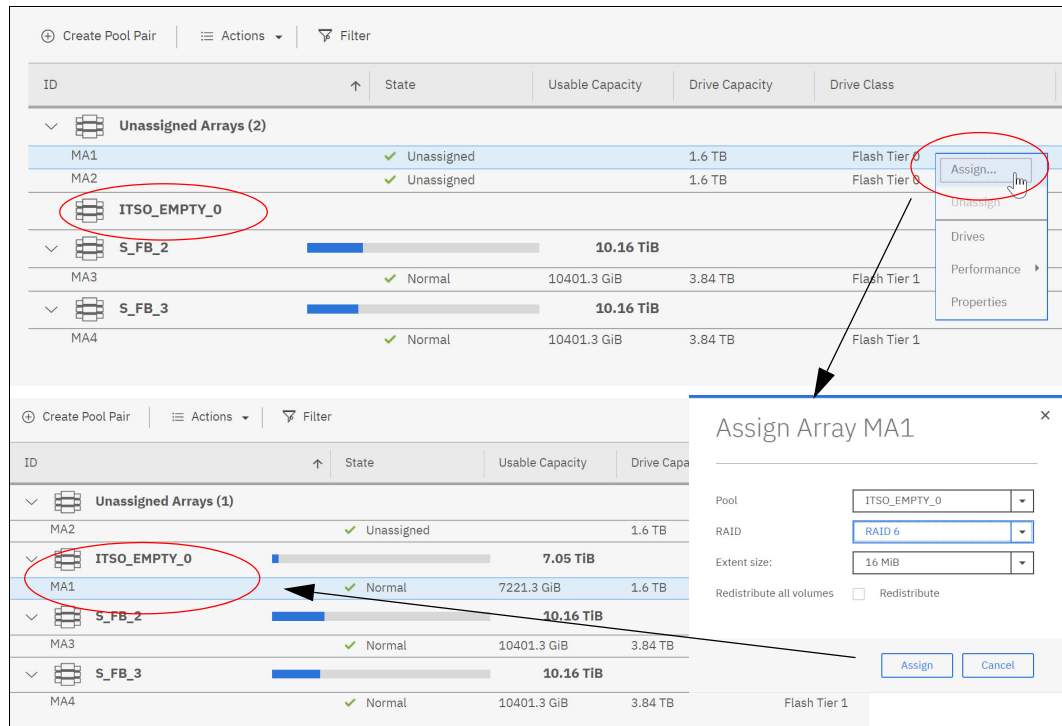


Figure 9-44 Assigning an array to a single pool

9.6.3 Creating FB volumes for open systems hosts

There are multiple paths to take when creating FB volumes. The most straightforward way is by using the **Create Volumes** tab on either the Volumes window or the Volumes by Pool window. You can also access this tab by clicking the **Actions** menu under Host or LSS. The maximum capacity for an FB volume is 16 TiB. The Storage Management GUI automatically distributes capacity across the two pools.

To create these FB volumes, complete the following steps:

1. From the system window, select the **Volumes** icon. Four options are provided, as shown in Figure 9-45 on page 279:
 - **Volumes** (All volumes are visible in single view.)
 - **Volumes by Pool** (Volumes are grouped by pool.)
 - **Volumes by Host** (Volumes are grouped by host.)
 - **Volumes by LSS** (Volumes are grouped by LSS.)

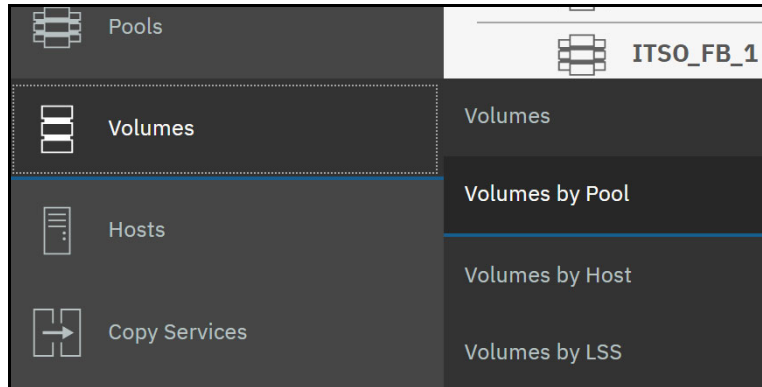


Figure 9-45 Views for Volumes

2. Selecting one of the first two options opens a view listing all the volumes or pools on the system. Figure 9-46 shows the Volumes by Pool view.

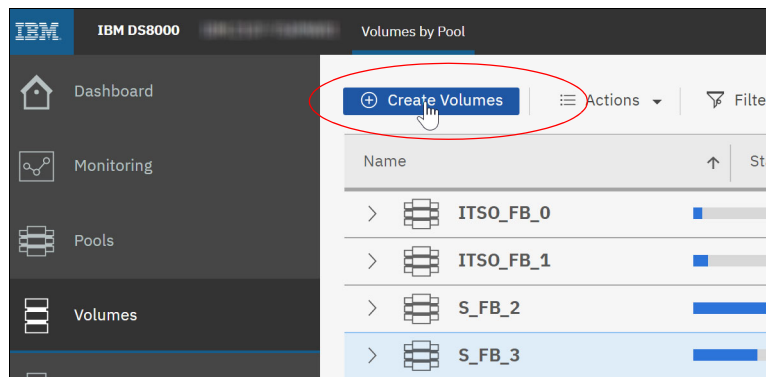


Figure 9-46 Volumes by Pool

3. From this view, click **Create Volumes**. The Create Volumes dialog opens, as shown in Figure 9-47.

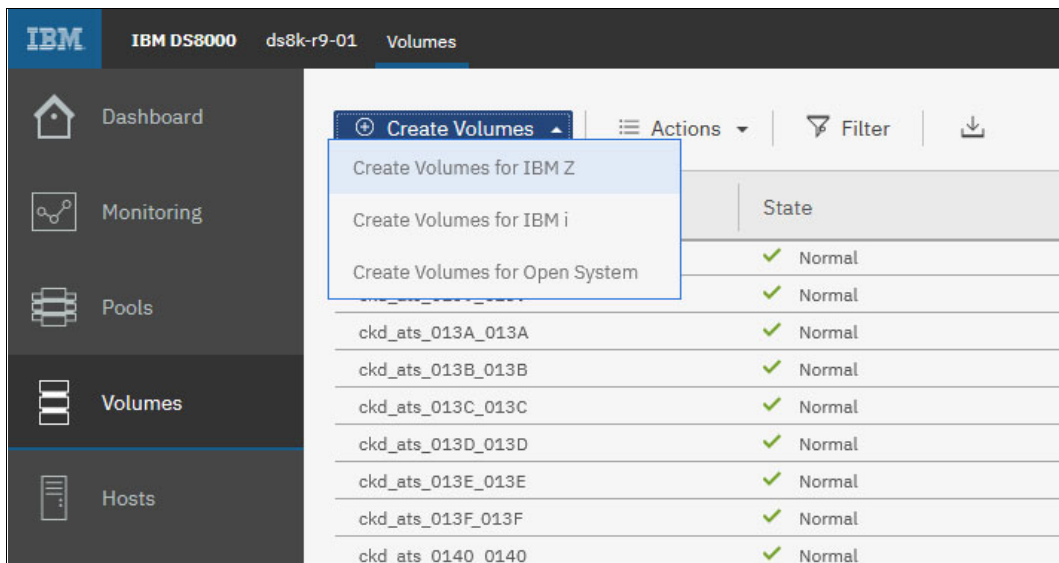


Figure 9-47 Create Volumes dialog

The DS GUI provides three *presets* or templates to create volumes:

- Create Volumes for IBM Z
Use this option to create CKD volumes for IBM Z host storage capacity provisioning.
- Create Volumes for IBM i
Use this option to create FB volumes for IBM i host storage capacity provisioning.
- Create Volumes for Open System
Use this option to create FB volumes for open system host storage capacity provisioning.

Creating FB volumes: Open systems

To create FB volumes for open systems hosts by using the open systems preset, complete these steps:

1. From the **Create Volumes** drop-down menu (Figure 9-47 on page 279), click **Create Volumes for Open Systems**.
2. The Create Volumes for Open Systems configuration dialog for open systems hosts opens (Figure 9-48). By default, the DS GUI tries to balance the volumes across the pool pair so that the workload is balanced across both nodes or central processor complexes (CPCs). So, by default it selects both CPC pools. An Administrator or Physical Operator user can select the pool that they want in the drop-down option at pool selection.

The screenshot shows a web-based configuration window titled "Create Volumes for Open Systems". At the top, there is a sub-header "New Volume Set" and a brief instruction: "Define one or more sets of volumes to be created. Each volume set is created identically across all LSSs in the range selected for that set." Below this is a toolbar with "New Volume Set", "Edit", and "Delete" icons. A table with columns "Volume names", "Pools", "LSSs", "Volume addresses", "Capacity", and "Provisioning" is shown, but it is empty with the text "No items found." The main configuration area includes:

- Pools:** Two dropdown menus for "Processor node 0" (selected: FB_2) and "Processor node 1" (selected: FB_3). Below each is a progress bar labeled "Usable 16.93 TiB".
- Host:** A dropdown menu labeled "Host or Cluster" with the text "-- Select target host or cluster --".
- Set:** Fields for "Name prefix" (ITSD_DS8900F_VOL), "Quantity" (2), "Capacity" (10 GiB), and "Provisioning" (Standard).

At the bottom, there are "Advanced", "Discard", and "Save" buttons.

Figure 9-48 Creating open systems (FB) volumes

3. Enter the following user-specified values:
 - Name prefix: User-defined name for volumes (a suffix ID sequence number is added during the creation process).
 - Quantity: Number of volumes to be created in selected pools.
 - Capacity: The capacity of the volumes to be created. Volumes can be configured in the following increments:
 - MiB, GiB, or TiB
 - Blocks

- (Optional) Host: Optionally, map the volumes to a target host or host cluster.
- Provisioning: Select type of Storage allocation:
 - Standard: Fully provisioned Volume
 - Thin provisioning: Thin provisioning defines logical volume sizes that are larger than the usable 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 (ESE) logical volumes.

The administrator or user, while creating the new volumes, can assign the address range to the volume in the Advanced section, as shown in Figure 9-49. It is possible to specify the volumes by using the T10 Data Integrity Field (DiF)/Protection Information. After you specify the volume set that you want to create, click **Save**. Then, you either create another one by selecting **⊕ New Volume Set**, or, once all the volume sets are specified, click **Create** to create them in a row all at once.

Figure 9-49 Using the Advanced section when creating volumes

Tips:

- ▶ By providing a target host or host cluster, you can create volumes and map them to the host in one step.
- ▶ Selecting the suitable range of addresses for the new volume set is important from the copy service planning point of view and the CPC preferred path affinity. After you create a volume, you cannot change its address.
- ▶ When FlashCopy is used on FB volumes, the source and the target volumes must have the same protection type, that is, they both must use T10-DiF or standard.

Creating volumes: IBM i

To create FB volumes for IBM i hosts by using the IBM i preset, complete the following steps:

1. Click the **Create Volumes for IBM i** option. The window that is shown in Figure 9-50 on page 282 opens.

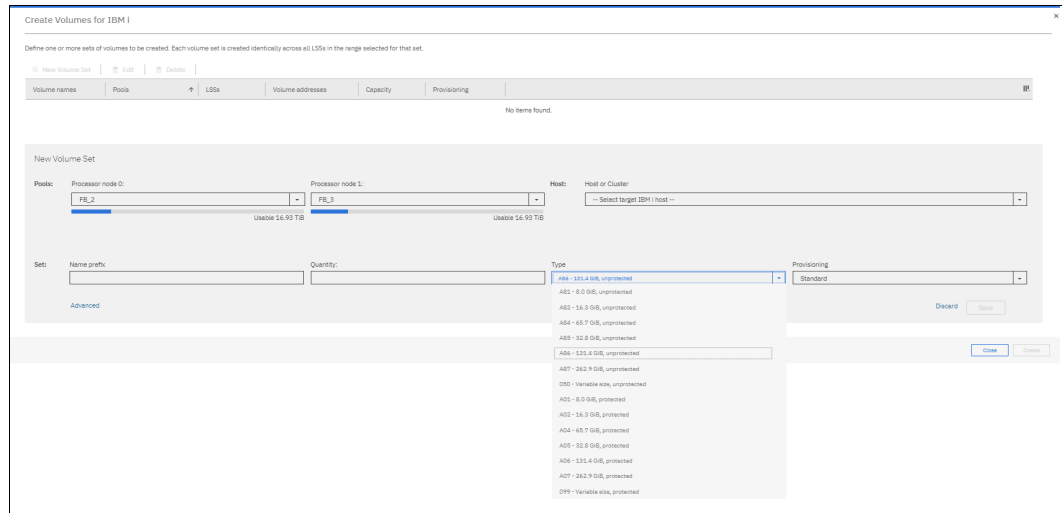


Figure 9-50 Creating IBM i volumes

2. Choose an FB pool from the available ones in each node.
3. Enter a Name prefix to identify the volumes.
4. Enter the Quantity of volumes to create.
5. Select the volume **Type** from a list of fixed capacities.
6. If variable type **050** (unprotected) or **099** (variable, protected) is selected, then the **Capacity** of the volumes (GiB or TiB) also must be specified.
7. In case you want volumes for Thin Provisioning (ESE), change the Provisioning field to Thin Provisioning (ESE).
8. As an option, you may specify 4-digit volume serial numbers for the volumes in the **Advanced** section. Use hexadecimal values (00-FE) for the LSSs.

Note: Release 9 and later supports Dynamic Volume Expansion (DVE) of IBM i 050 and 099 volume types in increments of 1 - 2000 GB. The minimum software level of the IBM i hosts must be IBM i 7.3 TR6 or IBM i 7.4 and later.

Optionally, you can map the volumes to a defined IBM i host in this step too.

Further volume sets can be prepared and saved before you create what is defined in a row.

9.6.4 Creating FB host attachments

To map FB volumes to open system hosts, complete the following steps:

1. Set the FC port topology.
2. Create open systems clusters (optional).
3. Create open systems hosts.
4. Assign host ports to hosts.
5. Assign FB volumes to open systems hosts.

Setting the Fibre Channel port topology

For an open system host to access FB volumes that are configured on the DS8900F, the host must be connected to the DS8900F through a FICON. The Fibre Channel Protocol (FCP) must be configured on the FC port so that the host can communicate with the volume on the DS8900F.

DS8900F has two kinds of host adapters: 4-port 16-gigabit Fibre Channel (GFC) and 4-port 32 GFC (referred to in the DS GUI as 16 Gbps or 32 Gbps). Each port can be independently configured to one of the following FC topologies:

- ▶ FCP: Also known as *FC-switched fabric* (which is also called *switched point-to-point*) for open system host attachment, and for Metro Mirror (MM), Global Copy (GC), Global Mirror (GM), and Metro/Global Mirror (MGM) connectivity
- ▶ FICON: To connect to IBM Z hosts, and for zGM connectivity

Note: With DS8900F, Fibre Channel Arbitrated Loop (FC-AL) is no longer supported.

To set the FC port topology for open system hosts, complete the following steps:

1. From the DS GUI left navigation menu, click **Settings** → **Network** and select **Fibre Channel Ports** to open the Fibre Channel Ports window (Figure 9-51).

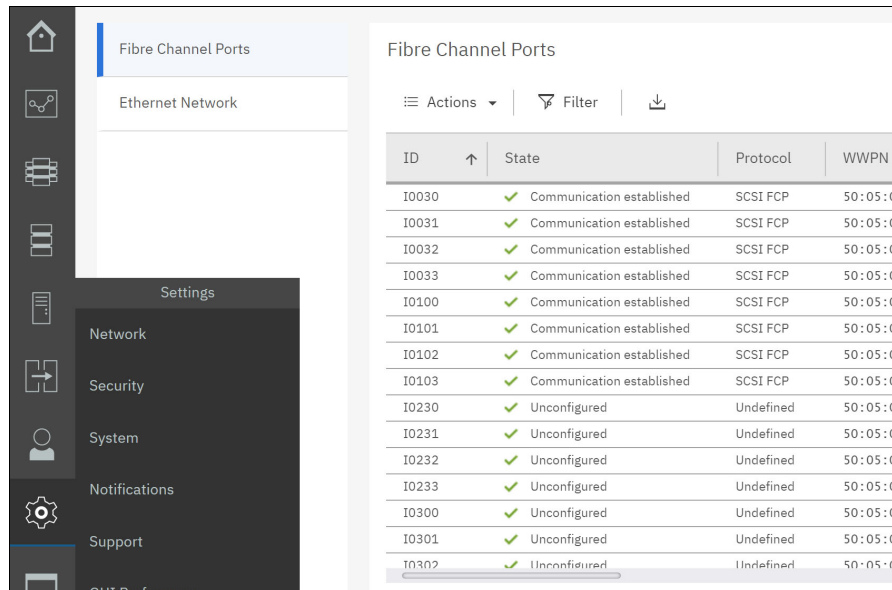


Figure 9-51 Fibre Channel Ports window

2. Select the port to modify. Multiple ports can be selected by using the Shift or Ctrl key.

- From the Actions tab, click **Modify Protocol** to open the Modify Protocol window, as shown in Figure 9-52.

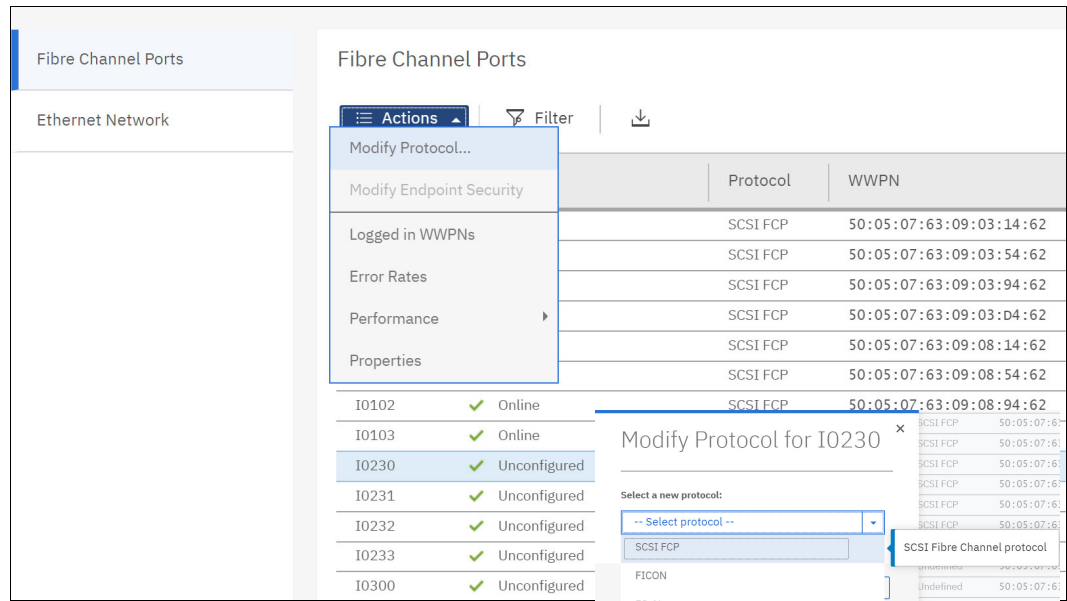


Figure 9-52 Modify Protocol window

- Choose from the available protocols to modify the selected host adapter port or ports. For open system hosts attachment, select **SCSI FCP** (Small Computer System Interface (SCSI) FCP).
- Click **Modify** to perform the action.

Creating open system clusters and hosts

To simplify and enable better management of volume-mapping operations to hosts, the DS8900F can group hosts of the same type into a host cluster. Hosts in a host cluster can have both shared and private volume mappings. This section describes how to configure host clusters and hosts by using the DS GUI.

For reference, a host port is the FC port of the host bus adapter (HBA) FC adapter that is installed on the host system. It connects to the FC port of the host adapter that is installed on the DS8900F.

To configure an open system cluster, complete the following steps:

- Create a cluster: Configure a cluster object to access the storage system.
- Create a host: Configure a host object to access the storage system.
- Assign hosts: Assign hosts to a cluster object.
- Assign a host port: Assign a host port to a host object by identifying one of the WWPNs of the HBA that is installed on the host system.
- Modify the FC port mask: Modify the FC port mask (on the DS8900F) to allow or disallow host communication to and from one or more ports on the system.

To configure an open system host, complete the following steps:

1. Create a host: Configure a host object to access the storage system.
2. Assign a host port: Assign a host port to a host object by identifying one of the WWPNs of the HBA that is installed on the host system.
3. Modify the FC port mask: Modify the FC port mask (on the DS8900F) to allow or disallow host communication to and from one or more ports on the system.

Creating clusters

To configure a cluster object, complete these steps:

1. Click the **Hosts** icon from the DS GUI navigation pane on the left.
2. Select **Hosts** from the menu, as shown in Figure 9-53.




Figure 9-53 Hosts menu in the navigation

3. The Hosts window opens, as shown in Figure 9-54. Click **Create Cluster**.
4. The Create Cluster window opens, as shown in Figure 9-54. Specify the name of the cluster, and click **Create**.

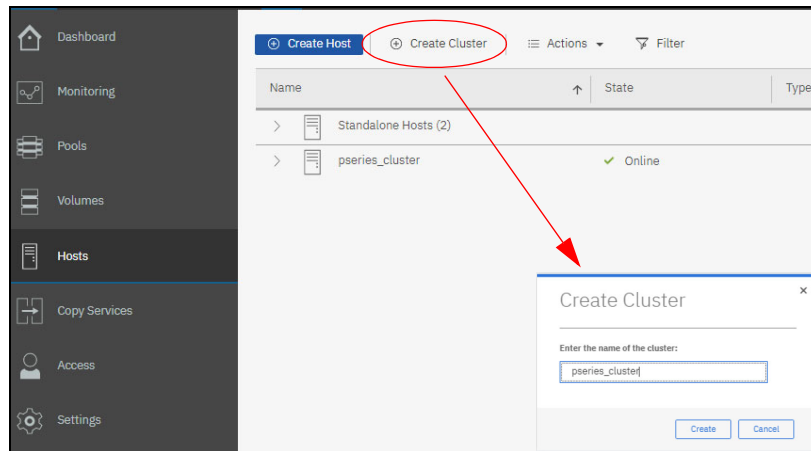


Figure 9-54 Hosts view: Create Cluster window

Creating hosts

To configure a host object, complete these steps:

1. Click the **Hosts** icon from the DS GUI navigation pane on the left.
2. Click **Hosts**, as shown in Figure 9-53.

- If any unassigned host ports are detected by the system, a suggested task window opens with two options, as shown in Figure 9-55.

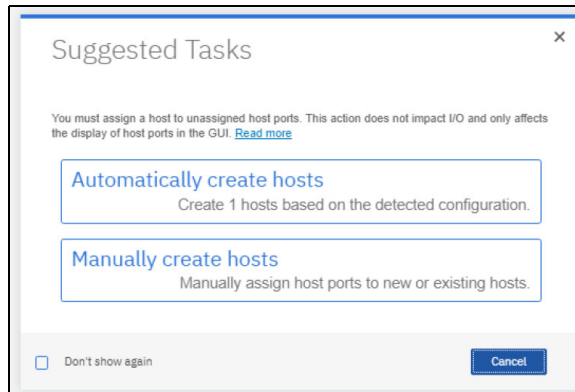


Figure 9-55 Hosts Suggested Tasks

Note: This window always appears when host port definitions are made by using the DS CLI (`mkhostport`) and are not yet fully reflected in the GUI. So, the GUI offers to move the CLI definitions fully into the GUI.

If canceled or closed, the Suggested Tasks window can be reopened by clicking the attention message that is shown in Figure 9-56.

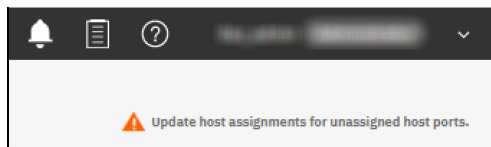


Figure 9-56 Updating host assignment attention

- Click **Automatically Create Hosts**. A list of detected hosts and unassigned host ports is displayed, as shown in Figure 9-57. Verify the list and click **Create** to complete the assignment task automatically.

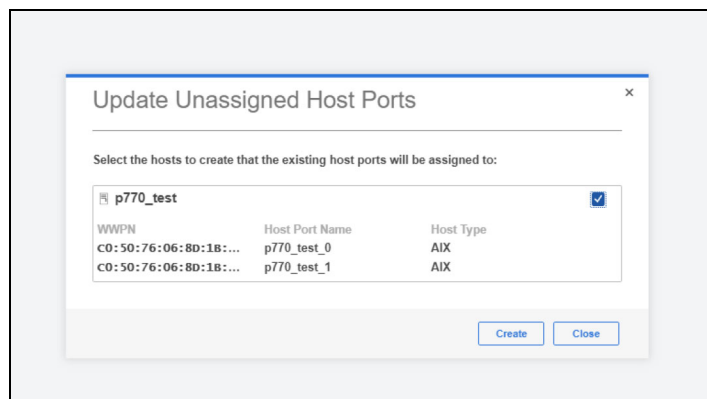


Figure 9-57 Update Unassigned Host Ports

- Select **Manually Create Hosts** to manually assign host ports to the existing hosts or click **Create Host** to create hosts to assign the ports.

6. In the Hosts window, click **Create Host**, as shown in Figure 9-58.

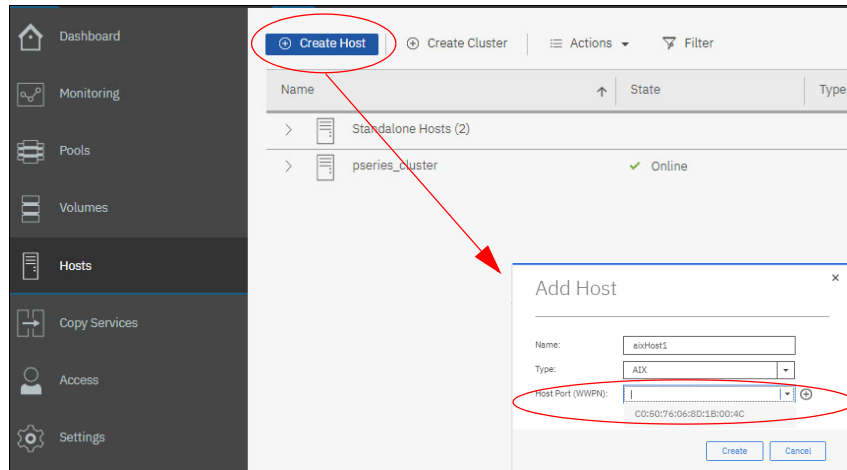


Figure 9-58 Create Host

7. The Add Hosts window opens (Figure 9-59). Specify the following items:

- Name: The user-defined name for the host to add.
- Type: The operating system (OS) of the host to add.
- Host port (WWPN): Optionally, provide the WWPN of the host port. If the host port logged in to the system, it can be selected from the Host Port (WWPN), list as shown in Figure 9-58.

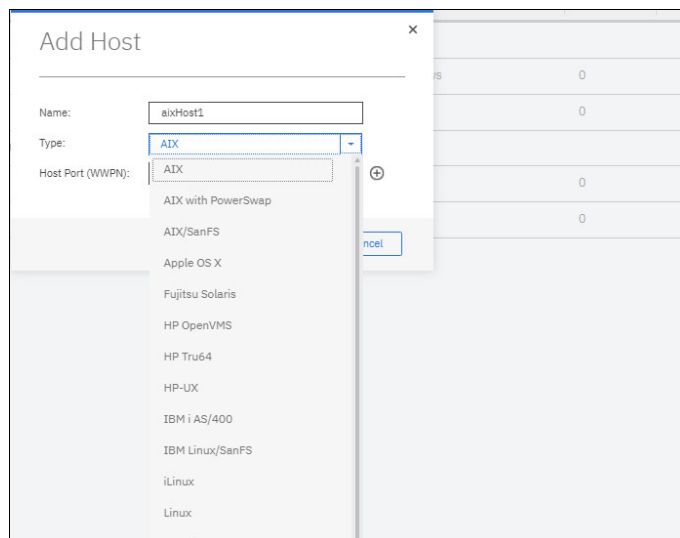


Figure 9-59 Add Host window that shows options for the host type

Assigning hosts to a cluster

If required, after a host is created, it can be assigned to a defined host cluster. To do so, complete these steps:

1. From the Hosts window, select the host to be assigned to a cluster.
2. Either right-click, or from the Actions tab, select **Assign to Cluster**, as shown in Figure 9-60.

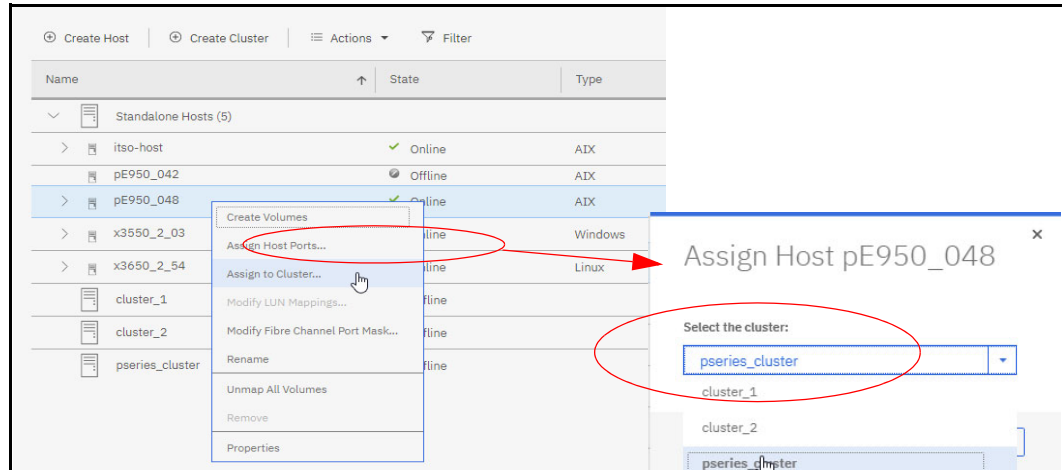


Figure 9-60 Assigning a host to a cluster

3. The Assign Host window opens. From the drop-down list, select the cluster to which to add the host.
4. Click **Assign** to complete the action.
5. Repeat the previous actions to add all hosts that are required in the cluster. After you complete all the hosts, assigned hosts are listed under the cluster in the Hosts window, as shown in Figure 9-61.

Name	State
Standalone Hosts (3)	
its0-host	Online
x3550_2_03	Online
x3650_2_54	Online
cluster_1	Offline
cluster_2	Offline
pseries_cluster	
pE950_042	Offline
pE950_048	Online

Figure 9-61 Hosts window showing the created cluster

Assigning host ports

After the host is added, host ports must be assigned to the defined host. To do so, complete these steps:

1. From the Hosts window, select the host to which to assign the host port. Either right-click, or from the Actions tab, select **Assign Host Port**, as shown in Figure 9-62.

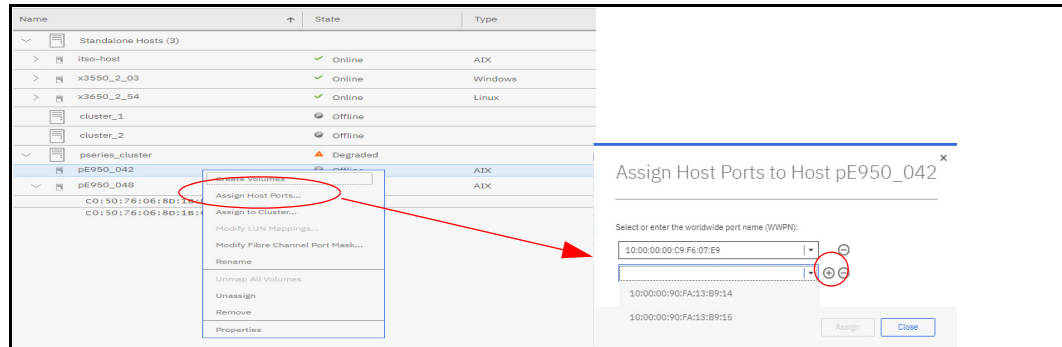


Figure 9-62 Assign Host Ports to Host

2. This action opens the Assign Host Ports to Host window.

If the host ports already logged in to the DS8900F (the host ports are zoned so that they are visible to the DS8900F), the available WWPNs are listed in the drop-down. If the host ports are not visible to the DS8900F, the WWPN of the host HBA must be added manually.

Select one of the WWPNs that are shown in the drop-down list, or manually enter the WWPN of the HBA of the host that you are adding. You can select multiple WWPNs by clicking +, as shown in Figure 9-62. Click **Assign** to complete.

Typically, most open system hosts have multiple FC connections to the DS8900F for redundancy and performance. Ensure that all additional host WWPNs that are connected for this host are defined to the host object by using the same procedure.

Note: When there are multiple FC connections to the DS8900F from a host, you should use native multipathing software that is provided by the host OS to manage these paths.

Modifying the Fibre Channel port mask

When the host is configured, by default it has access to all FC ports on the DS8900F system. To view them, select a host from the Hosts window, right-click, and then select **Properties**, as shown in Figure 9-63.

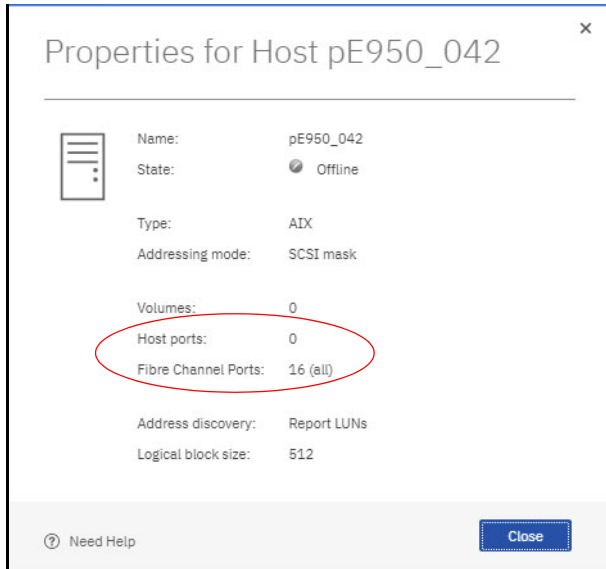


Figure 9-63 Host properties showing the Fibre Channel port mask

If the system administrator wants to restrict the FC ports that can communicate with the host, *FC port masking* must be defined. Modify the FC port mask to allow or disallow host communication to and from one or more ports on the system.

To define FC port masking, complete the following steps:

1. From the Hosts window, select the host to modify the FC port mask. Right-click and, or from the **Actions** tab, select **Modify Fibre Channel Port Mask**. A list of the DS8900F FC ports is displayed.
2. Select one or more ports to disallow. You can use the Ctrl or Shift key for multiple selections. Click **Save**, as shown in Figure 9-64.

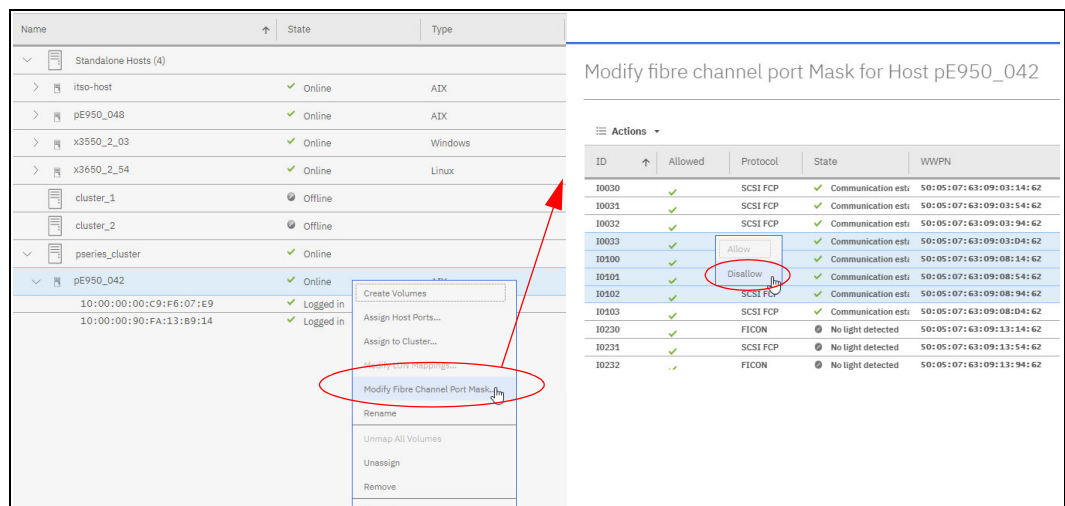


Figure 9-64 Modifying the Fibre Channel Port Mask

The properties of the selected host now reflect the number of FC ports that have access, as shown in Figure 9-65.

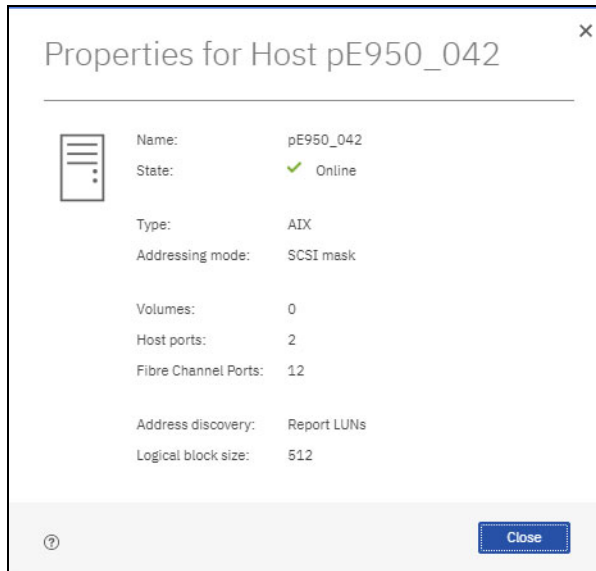


Figure 9-65 Host properties with FC port masking

9.6.5 Assigning FB volumes

FB volumes can accept I/O only from the host ports of hosts that are mapped to FB volumes. To map a volume to a host or a cluster, complete these steps:

1. From the DS GUI system window, select either the **Volumes** icon or **Hosts** icon, and then select **Volumes by Host**.
2. The Volumes by Hosts menu opens. Select the volumes to map from the **Unmapped Volumes** list, as shown in Figure 9-66. From the Actions tab, or by right-clicking, select **Map to Host or Cluster**.

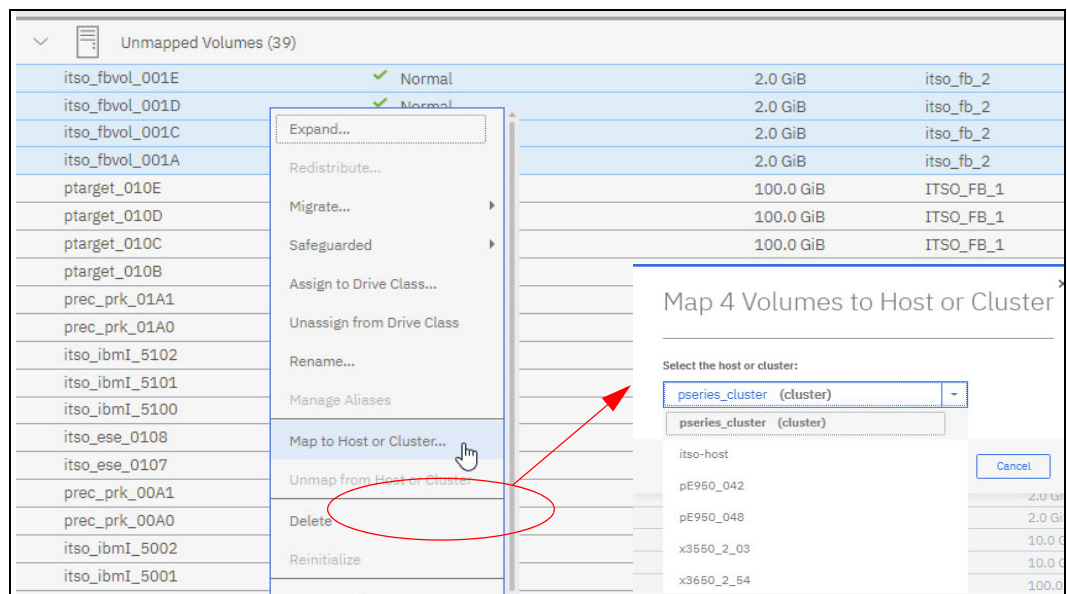


Figure 9-66 Mapping unmapped volumes to a host or cluster

- The Map Volume to Host or Cluster window opens. Select the host or cluster from the list of configured stand-alone hosts and clusters, as shown in Figure 9-66 on page 291, and then click **Map**.

Note: When mapping volumes to a cluster, volumes that are mapped to the cluster are public volumes that are seen by all hosts in the cluster. Volumes that are mapped to a single host in a cluster are private volumes.

It is the responsibility of the system administrator to ensure that the correct clustering software is implemented to ensure data integrity when a volume is mapped to more than one host.

Figure 9-67 shows a mixture of public and private volumes that are mapped to a cluster.

Name	State	Host Type	LUN	Mapping
> Unmapped Volumes (35)				
> Standalone Hosts (3)				
cluster_1	Offline			
cluster_2	Offline			
▼ pseries_cluster	Online			
▼ pE950_042	Online	AIX		
itso_fbvol_001A	Normal		4000401	Public
itso_fbvol_001C	Normal		4000401	Public
itso_fbvol_001D	Normal		4000401	Public
itso_fbvol_001E	Normal		4000401	Public
▼ pE950_048	Online	AIX		
_0080	Normal		4000408	Private
_0081	Normal		4000408	Private
_0180	Normal		4001408	Private
_0181	Normal		4001408	Private
itso_fbvol_001A	Normal		4000401	Public
itso_fbvol_001C	Normal		4000401	Public
itso_fbvol_001D	Normal		4000401	Public
itso_fbvol_001E	Normal		4000401	Public
pstress_0013	Normal		4000401	Private
pstress_0014	Normal		4000401	Private
pstress_0109	Normal		4001400	Private
pstress_010A	Normal		4001400	Private

Figure 9-67 Mapping volumes to a cluster

9.7 Logical configuration for Count Key Data volumes

This section describes the logical configuration for CKD volumes for IBM Z.

9.7.1 Configuration flow

As with the FB configuration, logical configuration for CKD volumes is simplified and can now be accomplished with a few steps:

- Create a CKD pool pair and assign arrays to the pools.
- Create the CKD LSSs.
- Create the CKD volumes.
- Configure the alias volumes (parallel access volumes (PAVs)).
- Configure the FC ports for FICON.

9.7.2 Creating CKD storage pools

For the best performance and a balanced workload, two pools must be created. The DS GUI helps the system administrator to create a balanced configuration by creating pools as a pair. The pools are configured so that one pool of the pair is managed by node 0, and the other pool of the pair is managed by node 1.

To create a CKD pool pair, complete these steps:

1. From the DS GUI window, select the **Pools** icon.
2. Click **Arrays by Pool** to open the Arrays by Pool window.
3. Click the **Create Pool Pair** tab.
4. In the Create Pool Pair window that is shown in Figure 9-68, select the number of arrays to assign to the pool pair.

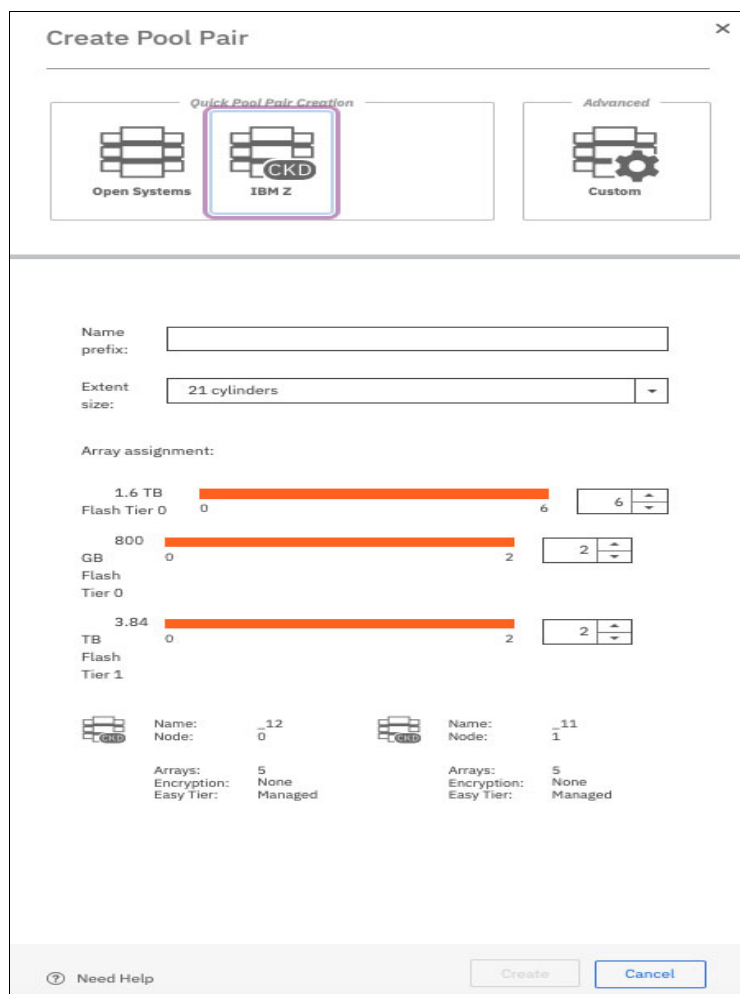


Figure 9-68 Creating the CKD pool pair and assigning arrays to the pool pair

5. If multiple drive classes are installed on the storage system, decide how many arrays of each drive class are required in each pool.
6. Ensure that storage type **CKD** is selected.
7. Assign a name to the pool pair. This name is used as the prefix for the pool pair ID.
8. Click **Create**.

Important: The CKD LSSs cannot be created in an address group that already contains FB LSSs. The address groups are identified by the first digit in the two-digit LSS ID.

- After the pool pair creation is complete, the arrays are assigned to the pool pair, as shown in Figure 9-69. The DS GUI configures the selected arrays for CKD storage and distributes them evenly between the two pools.

Note: You can automatically assign arrays when creating a pool pair. The arrays are created with the default RAID type RAID 6. To configure other supported RAID types, you must use the advanced configuration for pool creation, or assign the arrays manually to an existing storage pool from the unassigned arrays. For more information, see “Creating CKD Pools: Advanced configuration” on page 294.

ID	State	Usable Capacity	Drive Capacity	Drive Class
Unassigned Arrays (2)				
ITSO_CKD_0		8.10 KMod1		
MA2	✓ Normal	8099.0 Mod1	1.6 TB	Flash Tier 0
ITSO_CKD_1		8.10 KMod1		
MA1	✓ Normal	8099.0 Mod1	1.6 TB	Flash Tier 0

Figure 9-69 CKD pool pair that is created and arrays that are assigned to the pool pair

Creating CKD Pools: Advanced configuration

When creating CKD storage pools, you can also specify the extent size and RAID level. You can specify large 1,113 cylinder or small 21 cylinder extents for a storage pool or pool pair. The RAID level for the arrays can also be specified at creation time by using the Advanced (Custom) configuration option.

To create a custom CKD pool pair, complete these steps:

- From the System window, click the **Pools** icon.
- Click **Arrays by Pool** to open the Array by Pool window.
- Click the **Create Pool Pair** tab. The Create Pool Pair window opens.
- Select the **Custom** option.
- Specify the pool pair parameters:
 - Storage type: Ensure that **Count Key Data (CKD)** is selected.
 - Name prefix: Add the pool pair name prefix. A suffix ID sequence number is added during the creation process.
 - Extent size: Select **1,113 cylinders** for large extents or **21 cylinders** for small extents.
- Select from the listed drive types and select the number of arrays for each drive type that you want to assign to the pool pair.
- Choose the RAID level for the selected arrays.
- When pool pair parameters are correctly specified, click **Create** to proceed.

Figure 9-70 shows the custom configuration options for this procedure.

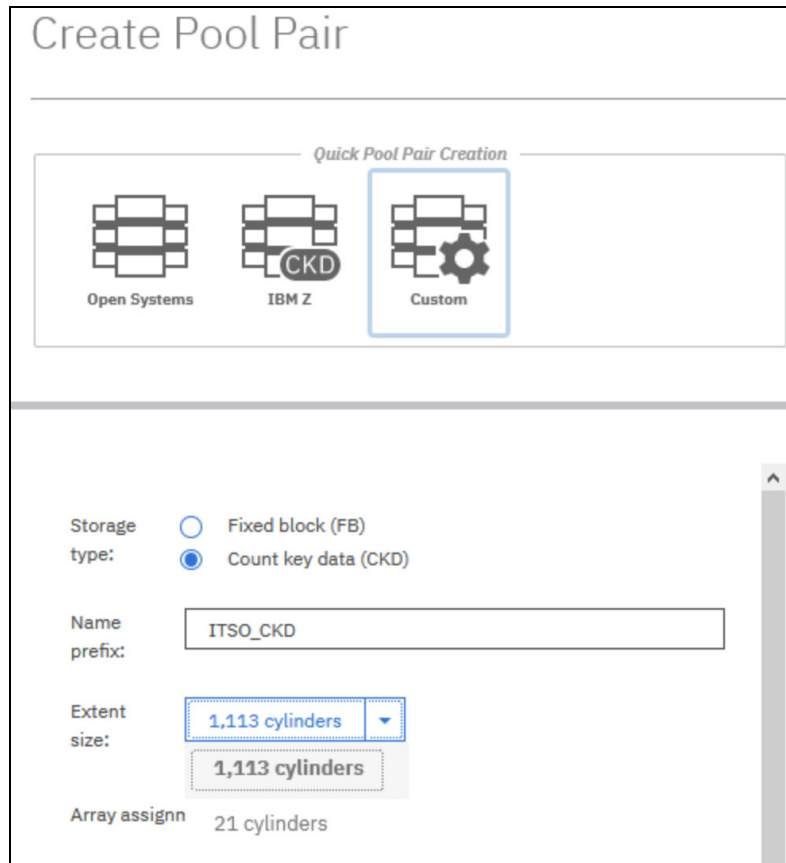


Figure 9-70 Configuring CKD storage pools with advanced options

Note: RAID 6 is the recommended and default RAID type for all drives.

RAID 5 is allowed for drives less than 1 TB with an accepted RPQ. When configuring RAID 5 for the supported drives, you must accept the disclaimer acknowledging that you understand the risks that are associated with RAID 5, as shown in Figure 9-71. A timestamped record with user information is created for audit purposes.

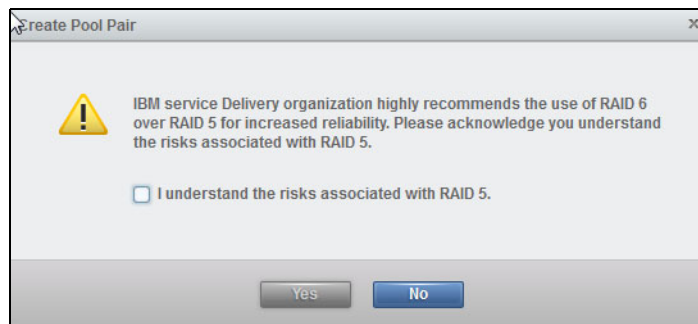


Figure 9-71 RAID 5 disclaimer for drives less than 1 TB

For more information about the supported drive types and available RAID levels for DS8900F models, see Chapter 2, “IBM DS8900F hardware components and architecture” on page 25.

Manually assigning arrays to existing pools

The system administrator can manually assign arrays, or unassign or reassign assigned arrays to existing pools when the current configuration must be modified, such as when the system administrator adds storage capacity. To do so, complete these steps:

1. From the Arrays by Pool window, select the array that you want to assign, then either right-click the array or select **Actions**, and then select **Assign**.
2. When you manually assign arrays, choose from an existing storage pool and define the RAID type.
3. Click **Assign**.

Figure 9-72 shows the Arrays by Pool window, which shows how to assign the arrays.

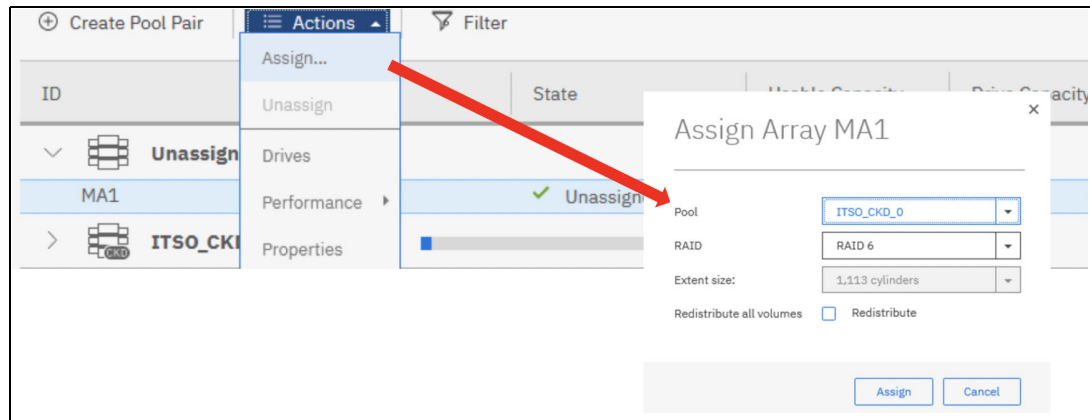


Figure 9-72 Manually assigning arrays to an existing pool

Creating a single pool

Occasionally, you are required to create only a single pool as opposed to creating a pool pair for balancing the workload. To create a single storage pool, see “Creating a single pool” on page 277.

9.7.3 Creating CKD logical subsystems

Note: You can create LSS ranges, exact volume address ranges, and aliases in one step. For an example, see 9.7.4, “Creating CKD volumes” on page 299.

The DS8000 LSS emulates a CKD storage control unit image (LCU). A CKD LSS must be created before CKD volumes can be associated to the LSS.

To create CKD LSSs, complete these steps:

1. Click the **Volumes** icon in the system window.
2. Select **Volumes by LSS** from the menu.

3. Click the **Create CKD LSSs** tab from the Volumes by LSS window, as shown in Figure 9-73.

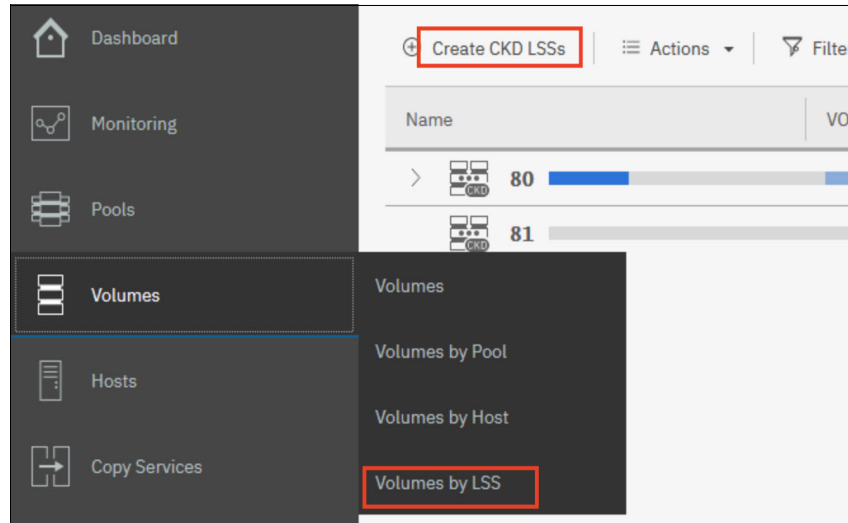


Figure 9-73 Creating CKD volumes by LSSs

4. The Create CKD LSSs window opens, as shown in Figure 9-74. Enter the required information. After you enter the values for the LSS range, subsystem identifier (SSID) prefix, and LSS type, click **Create**. The Need Help icon shows information about how the unique SSID for each LSS is determined based on the SSID prefix that is provided.

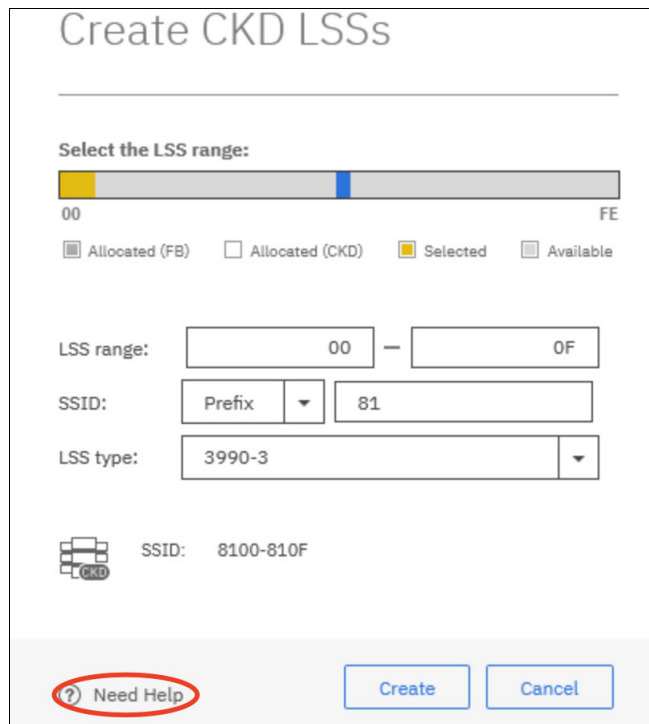


Figure 9-74 Defining CKD LSSs

Note: The CKD LSSs cannot be created in an address group that already contains FB LSSs. The address groups are identified by the first digit in the two-digit LSS ID.

Figure 9-75 shows the 16 LSSs that are created.

30	0 volume	SSID 3030
31	0 volume	SSID 3031
32	0 volume	SSID 3032
33	0 volume	SSID 3033
34	0 volume	SSID 3034
35	0 volume	SSID 3035
36	0 volume	SSID 3036
37	0 volume	SSID 3037
38	0 volume	SSID 3038
39	0 volume	SSID 3039
3A	0 volume	SSID 303A
3B	0 volume	SSID 303B

Figure 9-75 CKD LSSs that are created

- The unique SSID for each LSS is automatically determined by combining the SSID prefix with the ID of the LSS. The SSID can be modified if needed, as shown in Figure 9-76.

Important: This situation is important in an IBM Z environment where the SSIDs were previously defined in input/output definition files (IODFs) and might differ from the SSIDs that are automatically generated by the Storage Management GUI. Be careful when changing SSIDs because they must be unique in an IBM Z environment, and they are used in Copy Services definitions. A change must not be done unless you first removed all related copy services relationships (including PPRC paths).

Name	VOLSER	State	Capacity	SSID
80				SSID 8080
81				SSID 8081
82				SSID 8082
83				SSID 8083
84				SSID 8084
85	64 volumes	32 aliases		SSID 8085
86	64 volumes	32 aliases		SSID 8086
87	64 volumes	32 aliases		SSID 8087

Figure 9-76 Modifying the CKD LSS SSID

Note: Occasionally, the DS8900F GUI view does not immediately update after modifications are made. After you modify the SSID, if the view is not updated, refresh the GUI cache to reflect the change by clicking **Settings** → **Support** → **Troubleshooting** → **Refresh GUI cache**. For more information, see “Troubleshooting” on page 268.

9.7.4 Creating CKD volumes

To create CKD volumes, complete these steps:

1. Select the **Volumes** icon from the system window, and then select **Volumes** or **Volumes by Pool** from the menu.
2. Select the pools in which you want to create volumes.
3. Select the **Create Volumes for IBM Z** option under Create Volumes. The Create Volumes for IBM Z window opens, as shown in Figure 9-77.

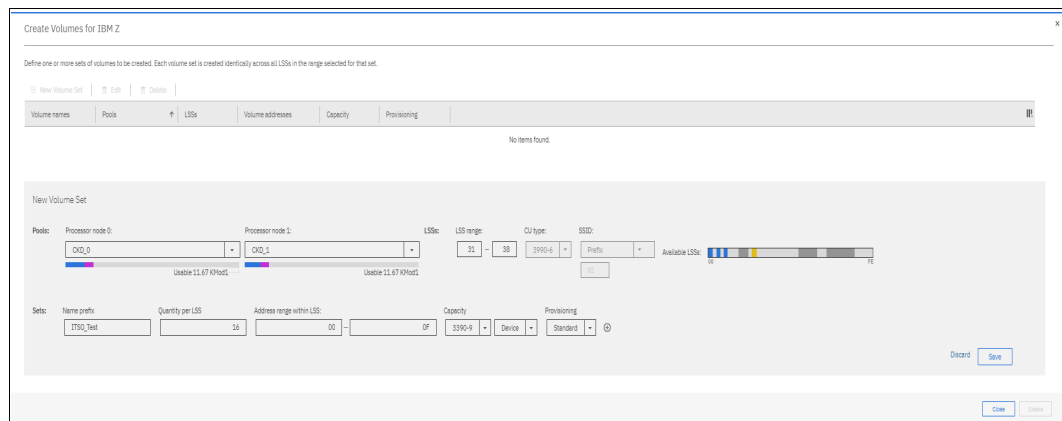


Figure 9-77 Creating multiple groups of CKD volumes

Note: The storage administrator can create configurations that specify new LSS ranges, exact volume address ranges, and aliases in one step.

4. Determine the LSS range for the volumes that you want to create.
5. Determine the name prefix and the quantity of volumes to create for each LSS.

Enter a prefix name and capacity for each group. The capacity can be specified in three ways:

- **Device:** Select one of these choices from the list: 3380-2, 3380-3, 3390-1, 3390-3, 3390-9, 3390-27, 3390-54, or 3390-A (extended address volume (EAV)). These device types have a fixed capacity that is based on the number of cylinders of each model. A 3390 disk volume contains 56,664 bytes for each track, 15 tracks for each cylinder, and 849,960 bytes for each cylinder. The most common 3390 model capacities are shown:
 - 3390-1 = 1113 cylinders
 - 3390-3 = 3339 cylinders
 - 3390-9 = 10017 cylinders
 - 3390-27 = 30051 cylinders
 - 3390-54 = 60102 cylinders

- Mod1: Emulates a 3390 Model 1. However, the number of cylinders is variable from 1 to the maximum available capacity in the pool.
 - Cylinders: Enter the number of cylinders for capacity based on a 3390 cylinder @ 849,960 bytes per cylinder.
6. Under **Provisioning**, select either **Standard** (thick volumes), or **Thin Provisioned (ESE)**.
 7. After completing the details, either define another volume set directly by using the ⊕ symbol, or click **Save** and validate the complete volume preset information. If you like to create another volume, click ⊕ **New Volume Set**. Otherwise, click **Create** to create all the volume definitions together in a row.

9.7.5 Creating CKD parallel access volumes

The DS8000 storage system supports the configuration and usage of PAVs. PAV is the concept of using multiple devices or aliases to address a single disk device. PAVs allow the definition of more unit control blocks (UCBs) to the same logical device, each using an extra alias address. For example, a direct access storage device at base address 1000 can have alias addresses of 1001, 1002, and 1003. Each of these alias addresses has its own UCB. Because there are four UCBs for a single device, four concurrent I/Os are possible. Writes to the same *extent* (an area of the disk that is assigned to one contiguous area of a file) are still serialized, but other reads and writes can occur simultaneously.

In the first version of PAV, the disk controller assigns a PAV to a UCB (static PAV). The second version of PAV processing, Workload Manager (WLM), reassigns a PAV to new UCBs from time to time (dynamic PAV).

The restriction for configuring PAVs is that the total number of base and alias addresses for each LSS cannot exceed 256 (00 - FF). These addresses must be defined in the IODF so that they match the correct type, base, or alias.

Typically, when you configure PAVs in the IODF, the base addresses start at 00 and increment toward FF. Alias addresses typically are configured to start at FF and decrement (decrease) toward 00. A system administrator might configure only 16 or 32 aliases for each LSS. However, no restrictions exist other than the total of 256 addresses that are available to the LSS (bases and aliases).

The DS GUI configures aliases in this manner, starting at FF and descending. The storage administrator can either configure many aliases against the LSS, in which case those aliases are assigned to the lowest address in the LSS. Alternatively, the system administrator can define any number of aliases to any specific base address. For more information about PAVs, see *IBM DS8900F and IBM Z Synergy DS8900F: Release 9.3 and z/OS 2.5*, REDP-5186.

Configuring parallel access volumes: Aliases

To configure aliases (PAVs), complete these steps:

1. Select the **Volumes** icon from the System window and then click **Volumes by LSS**.
2. Select the LSS for which to create aliases and right-click the LSS, as shown in Figure 9-78 on page 301.

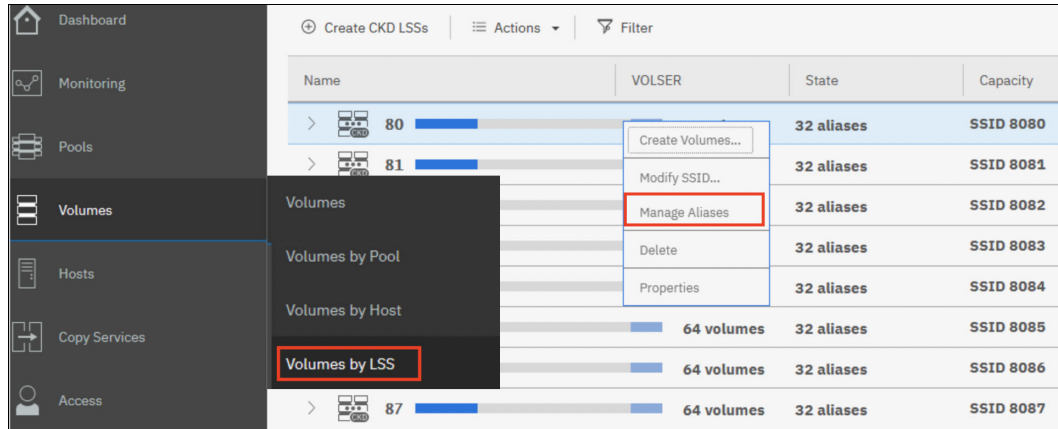


Figure 9-78 Creating aliases for the LSS

3. Select **Manage Aliases** to open the Aliases for LSS xx (where xx = 00 - FE) dialog box. Click **Create Aliases** to open the dialog box that is shown in Figure 9-79. Enter the number of aliases to create. The example in Figure 9-79 shows 32 aliases being created for LSS 80.

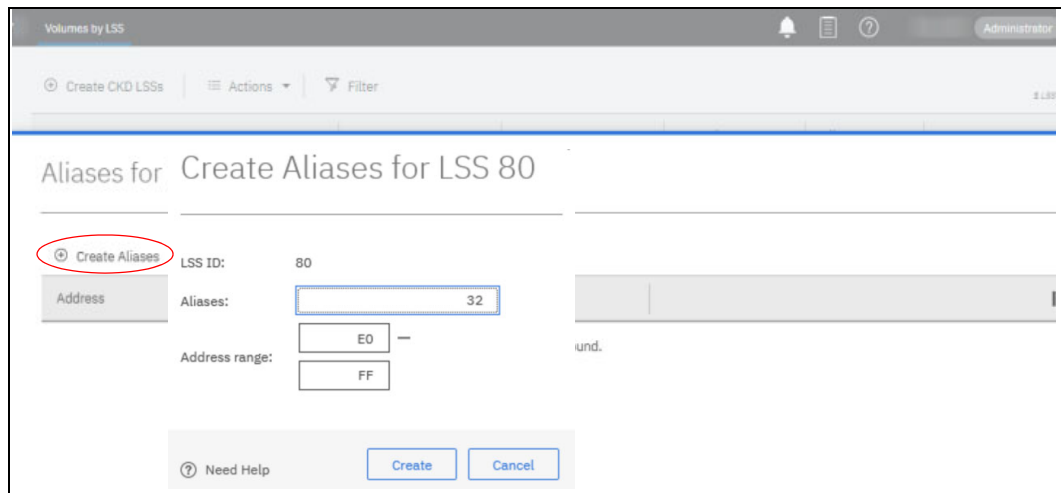


Figure 9-79 Creating 32 aliases for LSS 80

4. Click **Create**. The aliases are created for LSS 80, as shown in Figure 9-80.




Figure 9-80 Thirty-two aliases are created for LSS 80

- Click the arrow (>) icon to the left of the LSS to expand the LSS and display the base volumes that are assigned to the LSS.

Figure 9-81 shows the list of alias volumes for LSS 80.

The aliases are automatically created against the lowest base volume address in the LSS first. For example, in Figure 9-81, the 32 aliases are created against the lowest base volume address 8000 (ITSO_CKD_8000).

Name	VOLSER	State	Capacity	Aliases
80	64 volumes	32 aliases	SSID 8080	
ITSO_CKD_8000		✓ Normal	1.0 Mod1	32
ITSO_CKD_8001		✓ Normal	1.0 Mod1	0
ITSO_CKD_8002		✓ Normal	1.0 Mod1	0
ITSO_CKD_8003		✓ Normal	1.0 Mod1	0
ITSO_CKD_8004		✓ Normal	1.0 Mod1	0
ITSO_CKD_8005		✓ Normal	1.0 Mod1	0
ITSO_CKD_8006		✓ Normal	1.0 Mod1	0
ITSO_CKD_8007		✓ Normal	1.0 Mod1	0
ITSO_CKD_8008		✓ Normal	1.0 Mod1	0
ITSO_CKD_8009		✓ Normal	1.0 Mod1	0
ITSO_CKD_800A		✓ Normal	1.0 Mod1	0
ITSO_CKD_800B		✓ Normal	1.0 Mod1	0
ITSO_CKD_800C		✓ Normal	1.0 Mod1	0

Figure 9-81 Thirty-two aliases against the lowest base volume address

- To display the aliases, select the base volume with those aliases that are assigned to it and then click **Action** → **Manage Aliases**.

A list with the addresses of all aliases that are assigned to the base volume is displayed, as show in Figure 9-82.

Aliases for Volume ITSO_CKD__8000		
Address	Base Volume	LSS ID
80F5	ITSO_CKD_8000	80
80F6	ITSO_CKD_8000	80
80F7	ITSO_CKD_8000	80
80F8	ITSO_CKD_8000	80
80F9	ITSO_CKD_8000	80
80FA	ITSO_CKD_8000	80
80FB	ITSO_CKD_8000	80
80FC	ITSO_CKD_8000	80
80FD	ITSO_CKD_8000	80
80FE	ITSO_CKD_8000	80
80FF	ITSO_CKD_8000	80

Figure 9-82 List of aliases with their alias IDs starting at FF and in descending order

Note: The alias IDs start at FF and they are in descending order, as shown in Figure 9-83 on page 303.

- Aliases can also be created for a single base volume by selecting the base volume, right-clicking, and selecting **Action** → **Manage Aliases**. Then, select **Create Aliases**. Enter the number of aliases that you want for the base volume, as shown in Figure 9-83.

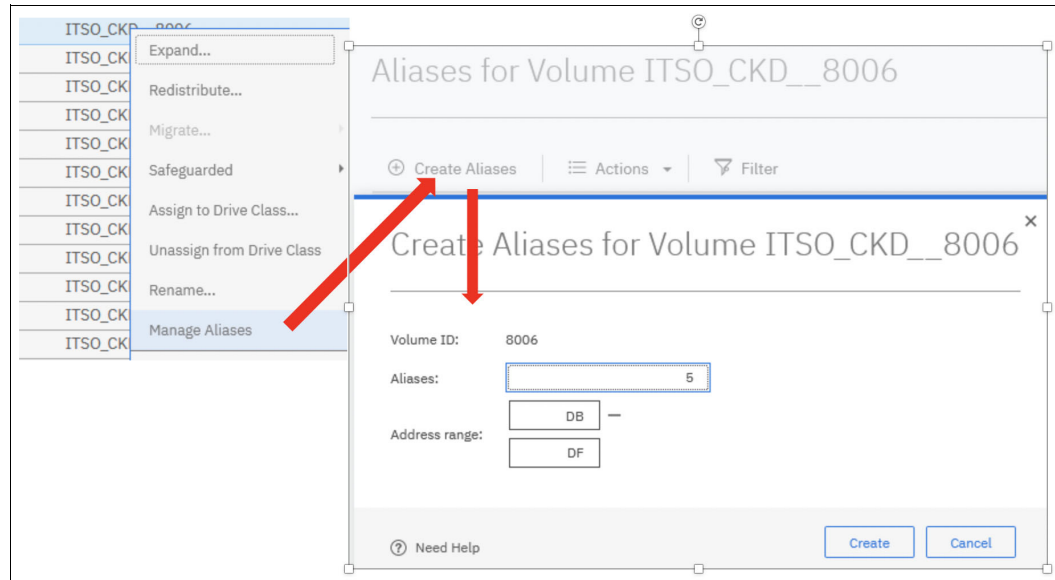


Figure 9-83 Configuring aliases for a single base volume

The five aliases for a single base volume address (ITSO_CKD_8006) are created with a starting address of DF and end with DB in descending order, as shown in Figure 9-84. (Alias E0-FF was created before).

Aliases for Volume ITSO_CKD__8006			
+ Create Aliases ☰ Actions 🔍 Filter			
Address	↑	Base Volume	LSS ID
80DB		ITSO_CKD__8006	80
80DC		ITSO_CKD__8006	80
80DD		ITSO_CKD__8006	80
80DE		ITSO_CKD__8006	80
80DF		ITSO_CKD__8006	80

Figure 9-84 List of five aliases that are created for a single base address

9.7.6 Setting the FC port protocols for IBM Z attachment

For an IBM Z host to access assigned CKD volumes, the host must connect to the DS8900F host adapter over FC. The protocol of the host adapter port must be set to FICON.

To set the FC port protocols of the FC ports that the host uses to communicate with the DS8900F, complete these steps:

1. Select **Settings** → **System** → **Fibre Channel Ports**, and then select **Actions** → **Modify Fibre Channel port protocols**.
2. Select one or multiple ports to modify and select **Actions** → **Modify**.
3. The Modify Protocol for the selected Ports window opens. Select the **FICON** protocol.
4. Click **Modify** to set the topology for the selected FC ports, as shown in Figure 9-85.

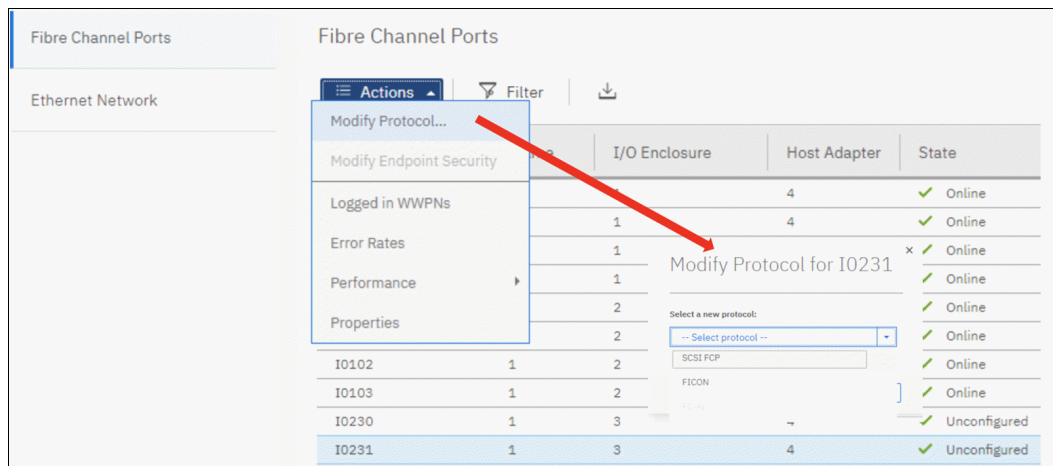


Figure 9-85 Modifying the FC port protocols

9.8 Expanding volumes

The DS8900F can dynamically expand a volume, which means that you can increase the capacity of a volume without taking the volume offline or without removing the existing data on the volume. This action is supported for both FB and CKD volumes. Additionally, this action is also supported for the IBM i variable size volume types 050 (unprotected) and 099 (protected).

The following example shows the steps that are required to expand an IBM i volume of type 099 (see Figure 9-86 on page 305):

1. Go to any Volumes view, such as **Volumes** → **Volumes by Pool**.
2. Select the volume and select **Actions** → **Expand**. You can also open the **Actions** menu by selecting the volume and right-clicking it.
3. The Expand Volume dialog opens. Enter the new capacity for the volume and click **Expand**.
4. A warning appears that informs you that certain OSs do not support this action, and it asks for confirmation to continue the action. Verify that the OS of the host to which the volume is mapped supports the operation, and click **Yes**.

A task window opens and is updated with progress on the task until it is completed.




Figure 9-86 Expanding a volume

The storage administrator can also expand the Safeguarded capacity, as shown in Figure 9-87. For more information, see *IBM Storage DS8000 Safeguarded Copy (Updated for DS8000 R9.2.3)*, REDP-5506.

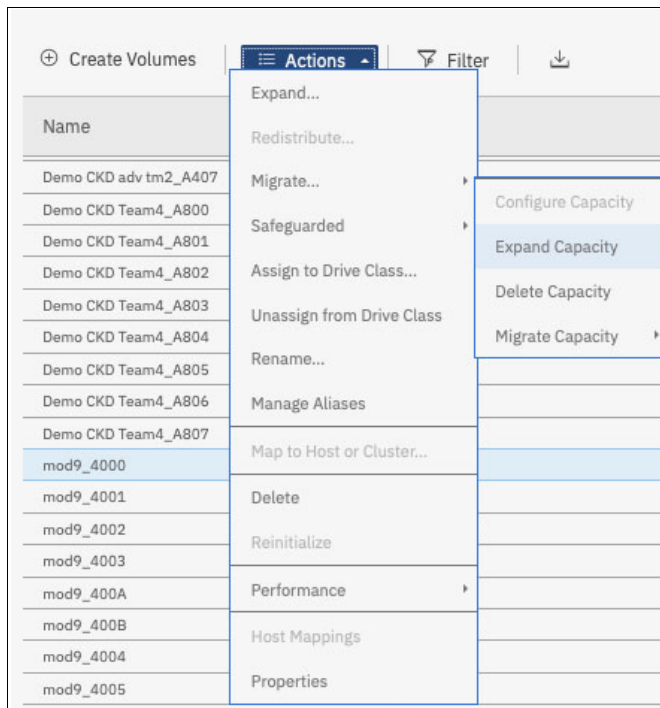


Figure 9-87 Expanding Safeguarded capacity

9.9 Deleting a pool

You can use the GUI to delete a pool that is no longer needed. If the pool contains volumes, you no longer need to first delete the volumes. If the volumes are not in use, it is now possible to delete a pool with volumes by entering the confirmation code that is presented by the GUI, as shown in Figure 9-88.

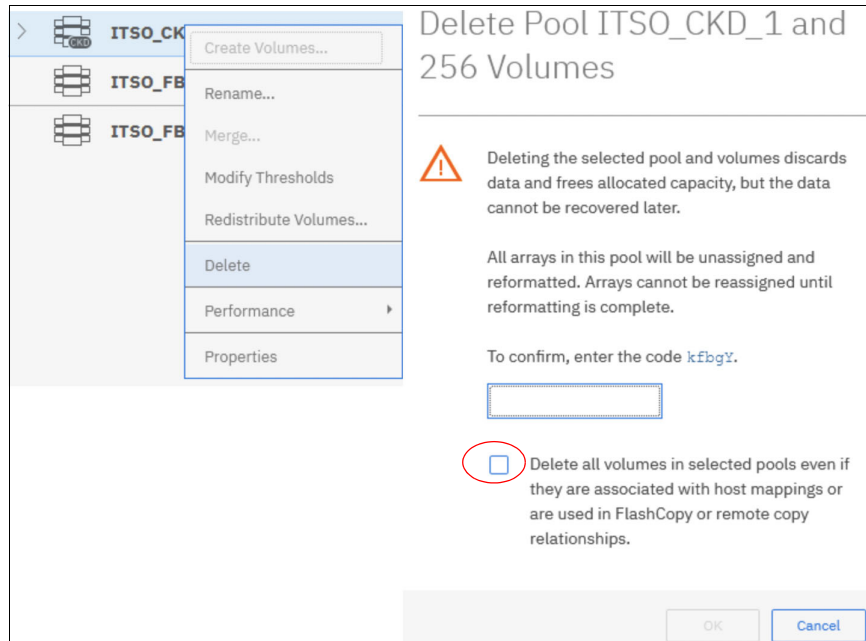


Figure 9-88 Deleting a pool with volumes

9.10 Deleting volumes

You can use the GUI to delete volumes that are no longer needed. By default, the GUI enforces safe deletion of volumes that are *in use*.

In the case of FB volumes, a volume is considered to be in use if it is in a CS relationship or if the volume received an I/O operation in the last 5 minutes. A CKD volume is considered in use if it participates in a CS relationship or if the IBM Z path mask indicates that the volume is in a grouped state or online to any host system.

You can instruct the GUI to force the deletion of volumes that are in use by selecting the optional checkbox in the **Delete Volumes** dialog box, as shown by #1 in Figure 9-89 on page 307. This setting does *not* apply to volumes that are in a Safeguarded Copy relationship.

The following example shows the steps that are needed to delete an FB volume that is in use:

1. Go to the hosts-centric Volumes view by selecting **Volumes** → **Volumes by Host**.
2. Select the volume to be deleted, and then select **Actions** → **Delete**, as shown by #2 in Figure 9-89 on page 307.

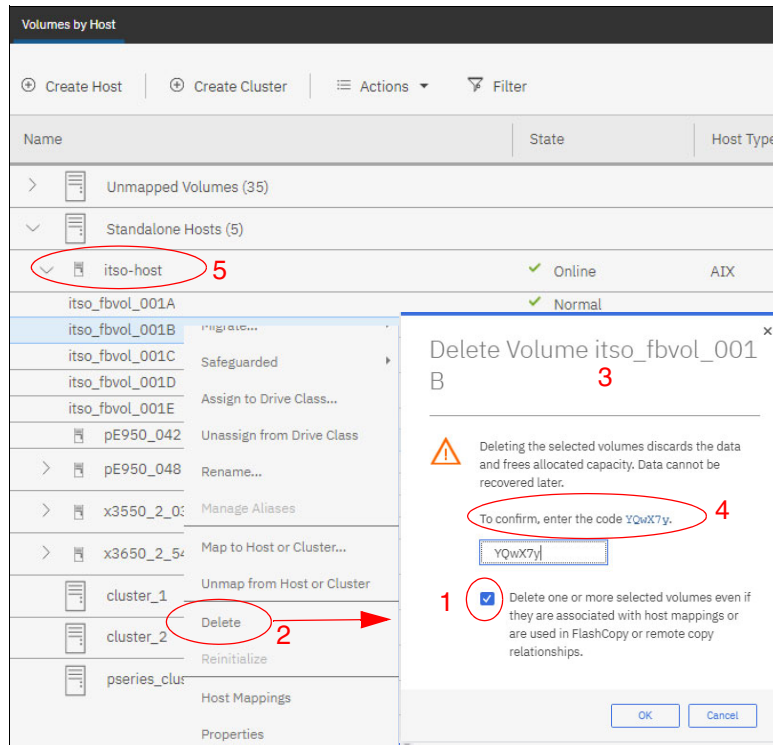


Figure 9-89 Deleting volumes

3. The Delete Volume window opens.
4. Confirm the action by entering the unique code that is displayed by the GUI into the window.
5. Because the selected volume is mapped to the host *itso-host* (#5), select the checkbox as shown by number 1 in the window and click **OK**.

A task window opens, and is updated with your progress on the task until it completes.

9.11 Reinitializing a thin-provisioned volume

Reinitializing a thin-provisioned volume frees storage, but also results in data loss. To prevent against accidental data loss, reinitialization of volumes requires more confirmation. To continue with the operation, confirm the action by entering the code that you receive. Additionally, for online volumes, select the **Force reinitialization** checkbox if you choose to bypass checking whether the selected volumes are mapped to a host or assigned to a path group, as shown in Figure 9-90.

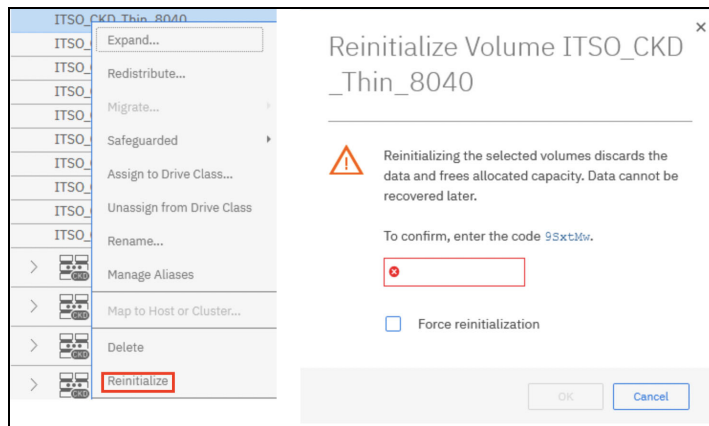


Figure 9-90 Reinitializing an online ESE volume

9.12 Easy Tier support

You can assign the volume to an Easy Tier drive class, or specify a tier that the volume should not use.

Assigning a volume to a drive class

Select a volume, and then click **Assign to Drive Class**, as shown in Figure 9-91 on page 309. A window with options that are based on the drive classes that are in the pool opens. If you select a drive class that is different from the drive class that is in use by the volume, Easy Tier moves the volume to the new drive class. The following options are available, depending on the drive classes that are in the pool:

- ▶ Flash Tier 0: Assign the volume to flash Tier 0.
- ▶ Flash Tier 1: Assign the volume to flash Tier 1.
- ▶ Flash Tier 2: Assign the volume to flash Tier 2.

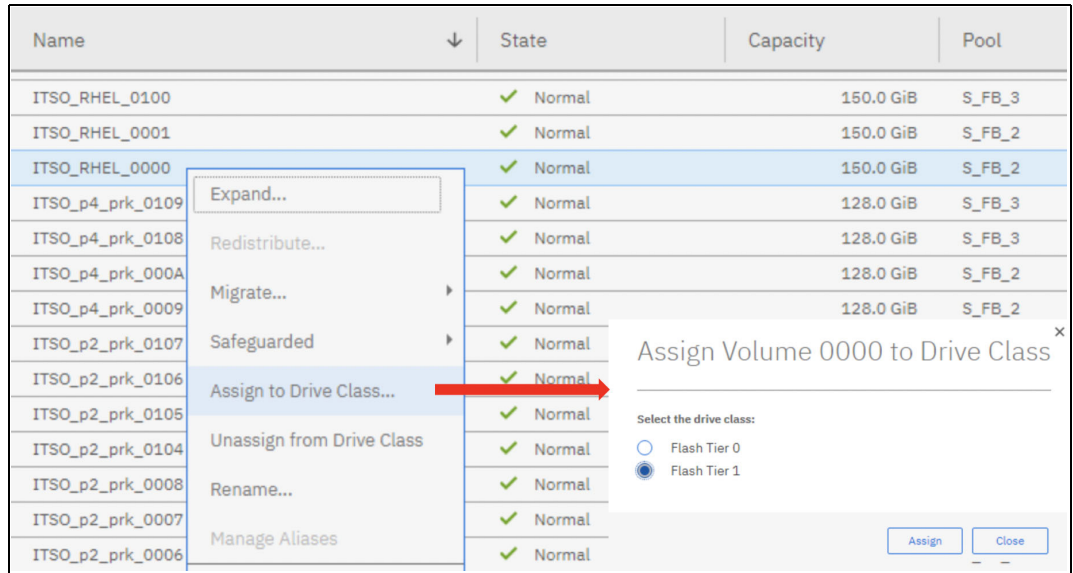


Figure 9-91 Assigning a volume to a drive class

Unassigning a volume from a drive class

By selecting the **Unassign from Drive Class** option, you can unassign a volume from an Easy Tier drive class tier and allow Easy Tier to manage the volume automatically again.

For more information about the settings that are available to configure Easy Tier, see “Easy Tier settings” on page 260.

For more information about Easy Tier, see *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667.

9.13 Monitoring system health

The DS8900F uses advanced predictive analysis to predict error conditions. The DS GUI provides tools to help monitor the health of the storage system in real time. If a system failure occurs, the system automatically provides notification. The Dashboard window in the DS GUI provides a visual representation of the DS8900F storage system. It displays system hardware states, and also provides access to actions that can be performed on the hardware components.

Figure 9-92 shows an example of the overall status of a system as displayed by the home dashboard.

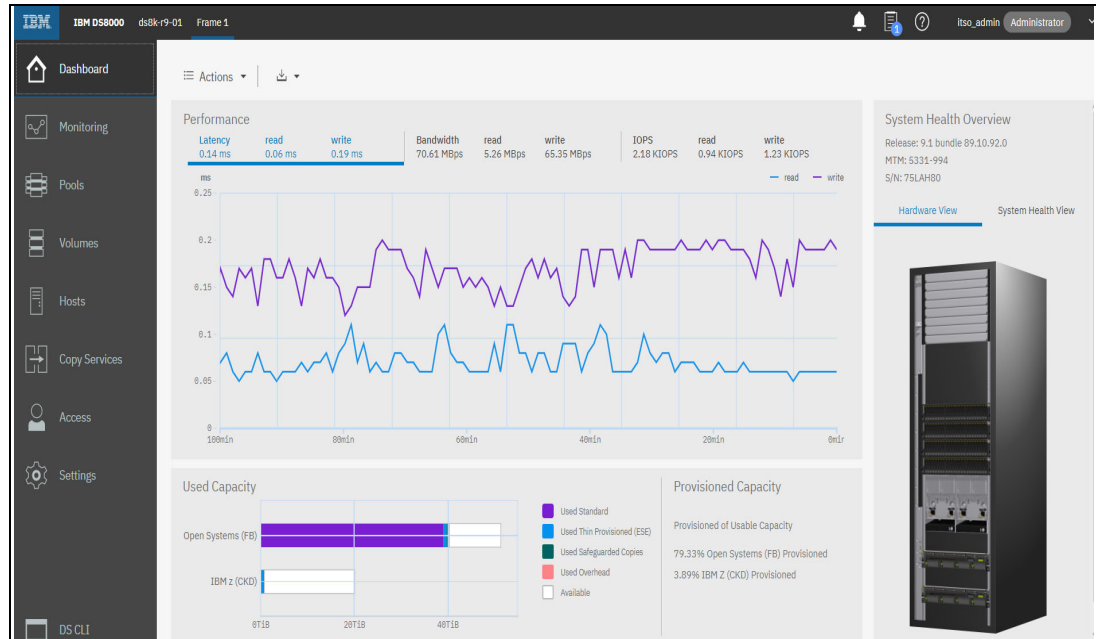


Figure 9-92 System Monitoring dashboard

The Dashboard window of the DS GUI includes the following elements:

- ▶ Performance

The performance of the system is displayed in real time, and it shows latency, bandwidth, and IOPS by default. Additionally, caching performance can be seen from the Performance tab at the top of the dashboard under **Actions/Performance**.

- ▶ Hardware resource alerts

The state of each hardware component of the storage system is displayed on the system window. Hardware components that need attention are highlighted.

- ▶ Used and Provisioned Capacity

Changes that affect the capacity of the storage system are displayed in the status area at the bottom of the Dashboard window.

- ▶ System Health Overview

This view lists the key hardware components of the system and their overall health. There are two views that are presented here:

- a. System health view

Changes that affect data accessibility and hardware states are displayed in the system health view on the right side of the Dashboard window. If a hardware error is displayed, you can also observe the hardware component that needs attention.

- b. Hardware view

A visual representation of the DS8900F system, which includes key hardware components. You can point to different system components for detailed information about that component. Click the component to see a more detailed view and show details and the state of that component. To restore the view of a component to its original position in the frame, click outside the frame.

► Alerting events

Error and warning alerts are displayed as badges on the Alerts (Bell) icon in the banner of the DS GUI (shown in Figure 9-93). Click the **Alerts** icon to see the alerts. Click the specific alert to view the corresponding event in the Events window.




Figure 9-93 Dashboard Events

9.13.1 Hardware components: Status and attributes

The DS GUI can identify the different hardware components of the storage system through a Hardware view, and it offers a detailed System Health overview of the system. Five of these hardware components are listed as an example:

- Processor nodes
- HMCs for management
- Storage enclosures with drives
- I/O enclosures with host adapters and DAs
- FC ports

This section provides more information about these hardware components, including how to see more information about them from the system window.

Note: For the DS8910F Rack-Mounted model 993, the status of the various hardware components is available from the System Health Overview. You can click each component to see more details about it. There is no hardware view for this model.

Processor nodes

Two processor nodes exist that are named ID 0 and ID 1. Each node consists of a CPC and the Licensed Internal Code (LIC) that runs on it. You also can display the system health overview by clicking the **System Health View** icon, as shown in Figure 9-94.

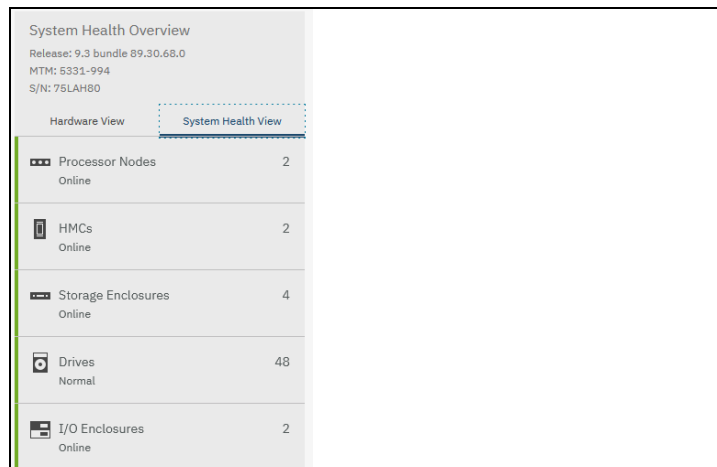


Figure 9-94 System Health Overview

For more information about the processor nodes, click **Processor Nodes**.

ID	State	Release	Processor	Memory	Location Code
1	Online	R9.2 bundle 89.20.131.0	POWER9 (22-way)	2176 GB	U78D2.001.WZS0TEZ
0	Online	R9.2 bundle 89.20.131.0	POWER9 (22-way)	2176 GB	U78D2.001.WZS0TF1

Figure 9-95 Processor Nodes detailed view

Here are the node attributes that are shown in Figure 9-95:

- ▶ ID: The node identifier, which is node 0 or node 1.
- ▶ State: The current state of the node is shown:
 - Online: The node is operating.
 - Initializing: The node is starting or not yet operational.
 - Service required: The node is online, but it requires service. A call home was initiated to IBM Hardware Support.
 - Service in progress: The node is being serviced.
 - Drive service required: One or more drives that are online require service. A call home was initiated.
 - Offline: The node is offline and non-operational. A call home was initiated.
- ▶ Release: The version of the Licensed Machine Code (LMC) or hardware bundle that is on the node.
- ▶ Processor: The type and configuration of the processor that is on the node.
- ▶ Memory: The amount of raw system memory that is installed in the node.
- ▶ Location Code: Logical location of the processor node.

Hardware Management Console

The DS8900F has two HMCs. The HMC provides a standard interface on a dedicated console to control managed systems. Point to the HMC component (Figure 9-94 on page 311) and click **HMCs** to display the detailed attributes for both HMCs, as shown in Figure 9-96.

Name	State	Role	Release	Management IP	Disk	Memory	Machine Model	Machine Type	Location Code
ds8k-r9-11...	Online	Secondary(2)	R9.2 bundle 89.20.131.0		256GB	32GB	L47	1700	U1700L47.1200927
ds8k-r9-11...	Online	Primary(1)	R9.2 bundle 89.20.131.0		256GB	32GB	L47	1700	U1700L47.1200824

Figure 9-96 HMCs detailed view

Here are the attributes that are displayed for the HMC component in Figure 9-96 on page 312:

- ▶ Name: The name of the HMC as defined by the user.
- ▶ State: The status of the HMC is shown:
 - Online: The HMC is operating normally.
 - Code updating: The HMC software is being updated.
 - Service required: The HMC is online, but it requires service. A call home was initiated to IBM Hardware Support.
 - Offline with redundancy: The HMC redundancy is compromised. A call home was initiated to IBM Hardware Support.
 - Offline: The HMC is offline and non-operational.
- ▶ Release: The version of the LMC that is installed on the HMC.
- ▶ Host address: The IP address for the host system of the HMC.
- ▶ Role: The primary or secondary HMC.
- ▶ Location Code: The logical location code for the HMC. If the HMC is external, the location is identified as off-system.

Storage enclosures

A *storage enclosure* is a specialized chassis that houses and powers the flash drives in the DS8900F storage system. The storage enclosure also provides the mechanism to allow the drives to communicate with one or more host systems. All enclosures in the DS8900F are High-Performance Flash Enclosures (HPFEs). These enclosures contain flash drives, which are Peripheral Component Interconnect Express (PCIe)-connected to the I/O enclosures.

To view detailed information about the enclosures that are installed in the system, select **Storage Enclosures**, as shown in Figure 9-94 on page 311. The attributes of the storage enclosure are shown in Figure 9-97.

ID	Frame	State	Drive Class	Drive Capacity	Installed	Location Code
T8	1	Online	Flash Tier 2	7680 GB	04/28/2021 01:17:07 PM	U2107.D04.J0CL004
T6	1	Online	Flash Tier 0	1600 GB	04/28/2021 01:17:07 PM	U2107.D04.J12K036
T7	1	Online	Flash Tier 2	7680 GB	04/28/2021 01:17:07 PM	U2107.D04.J04X018
T5	1	Online	Flash Tier 0	1600 GB	04/28/2021 01:17:07 PM	U2107.D04.J126012

Figure 9-97 Storage Enclosures detailed view

- ▶ ID: The storage enclosure number.
- ▶ State: The current state of the storage enclosure is shown:
 - Online: The storage enclosure is operating normally.
 - Service requested: A service request to IBM was generated for one or more drives within the storage enclosure.
 - Service required: The storage enclosure is online, but it requires service.

- Service in progress: The storage enclosure is being repaired.
- Offline: The storage enclosure requires service.
- ▶ Drive capacity: The raw capacity of the flash drives that are installed in the storage enclosure.
- ▶ Drive class: The drive class (tier) of the drives in the storage enclosure. Examples include Flash Tier 0, Flash Tier 1, and Flash Tier 2 drives.
- ▶ Installed: The time and date when the enclosure was installed.
- ▶ Location Code: The logical location code, which includes the serial number of the storage enclosure.

A *drive* is a data storage device. From the GUI perspective, a drive can be either a Flash Tier 0, Flash Tier 1, or Flash Tier 2 drive. To see the data storage devices and their attributes from the Hardware view, click a storage enclosure when the magnifying glass pointer appears (Figure 9-92 on page 310). This action shows the storage enclosure and installed storage devices in more detail (Figure 9-98). You can also select **Drives** from the System Health Overview to display information for all installed drives.



Figure 9-98 Drive detail view in a storage enclosure

The attributes for the drives in Figure 9-98 are described:

- ▶ ID: The drive, storage enclosure, and frame ID
- ▶ State: The current state of the drives is described:
 - Normal: Operational.
 - Initializing: The drive is being prepared for operation.
 - Service: The drive is being serviced.
 - Offline: The drive requires service.
- ▶ Capacity: The raw capacity of the drive
- ▶ Class: The drive class

I/O enclosures

The *I/O enclosure* contains the I/O adapters. To see the I/O adapters and their attributes from the Hardware view, click an I/O enclosure (Figure 9-92 on page 310) when the magnifying glass pointer appears.

This action displays the I/O enclosure adapter view (rear of the enclosure).

To see the attributes for all installed DAs or host adapters in the System Health overview, select **Device Adapters** or **Host Adapters** from System Health Overview (Figure 9-94 on page 311) to open the dialog that is shown in Figure 9-99.

ID	Frame	State	Host Adapters	Device Adapters	Location Code
4	1	Online	3	2	U1500.1B4.RJAB0BK
3	1	Online	3	2	U1500.1B3.RJAB0BP

Figure 9-99 I/O Enclosures

The attributes for the I/O enclosure are described in the following list:

- ▶ ID: The enclosure ID.
- ▶ State: The current state of the I/O enclosure is one of the following states:
 - Online: Operational and normal.
 - Offline: Service is required. A service request to IBM was generated.
 - Service: The enclosure is being serviced.
- ▶ Location Code: Logical location code of the I/O enclosure.
- ▶ DA: Number of DAs that are installed.
- ▶ Host adapter: Number of host adapters that are installed.

The *FC ports* are ports on a host adapter that connect the DS8900F to hosts, switches, or another storage system either directly or through a switch.

In the Hardware view, point to a port in the enclosure, and detailed information about the port appears. The port topology can be modified by selecting the port, right-clicking, and selecting **Modify Fibre Channel Port Protocol**, as shown in Figure 9-100.

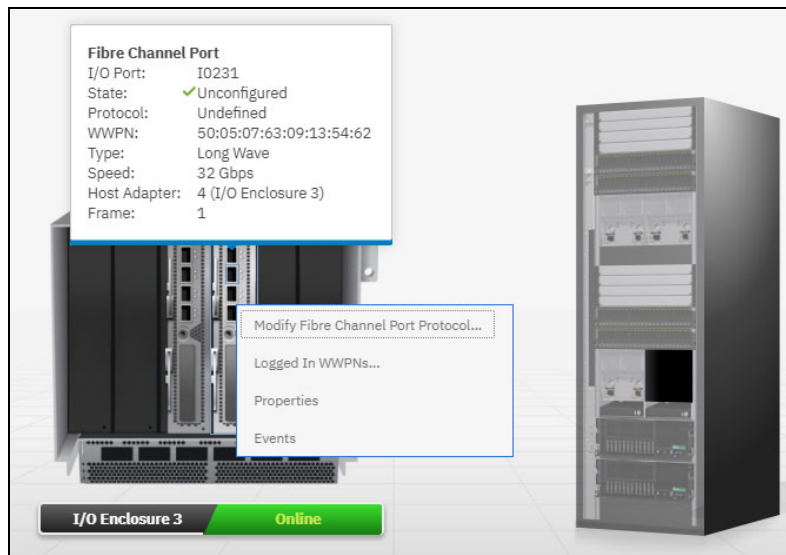


Figure 9-100 I/O enclosure showing host adapter and FC ports

Here are the attributes for the FC ports:

- ▶ ID: The identification number of the FC port.
- ▶ State: The current state of the FC port is one of the following states:
 - Communication Established: The FC port is operating normally.
 - No Light Detected: There is no light that is detected on the port, so the port cannot process I/O.
 - Unconfigured: The FC port protocol is not configured.
 - Offline: The FC port requires service.
 - Protocol: Use one of these FCPs:
 - SCSI FCP: A method for transferring SCSI commands with FCP.
 - FICON: FICON is a high-speed I/O interface for IBM Z host connections.
 - WWPN: The unique 16-digit hex number that represents the WWPN of the FC port.

9.13.2 Viewing components health and state from the system views

The Hardware view and System Health overview are useful tools in the DS GUI to visualize the state and health of hardware components in the system. Two sample scenarios are illustrated here:

- ▶ **Failed drives:** Figure 9-101 shows two failed drives in one of the storage enclosures of the system. Clicking this enclosure from the Hardware view provides a detailed view of the enclosure with all the drives, including the failed drives. Hovering your cursor over the failed drives provide more information about them.




Figure 9-101 Hardware view with failed drives

The same information can be obtained by clicking the **Show System Health Overview** icon, as shown in Figure 9-101.

The System Health Overview opens and shows the failed state of the drives, as shown in Figure 9-102.

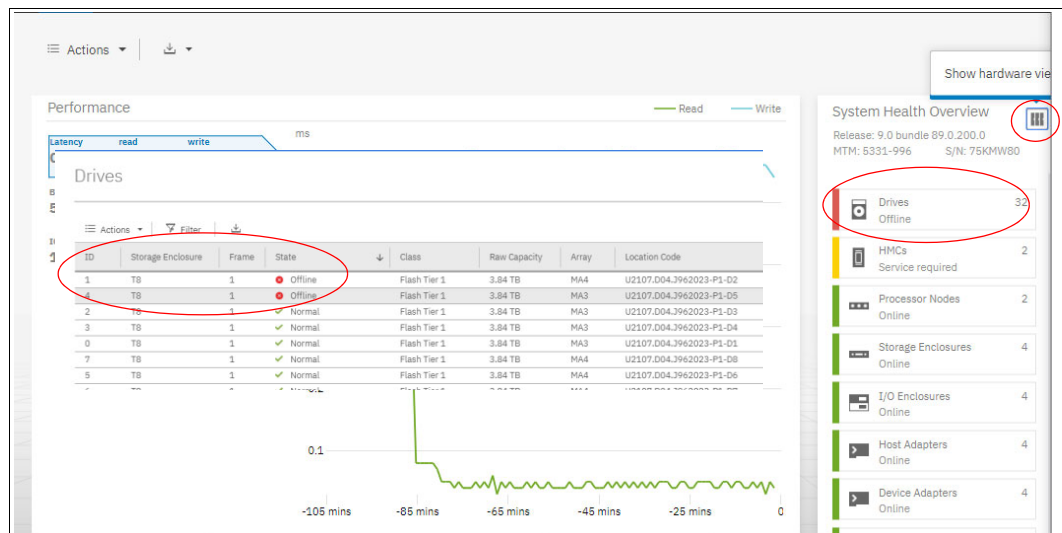


Figure 9-102 System Health Overview with failed drives

- Failed processor node: Figure 9-103 shows the Hardware view with a failed processor node. Hovering your cursor over the node provides more information about this component.




Figure 9-103 Hardware view with a failed processor node

The same information can be obtained from the System Health Overview, as shown in Figure 9-104.

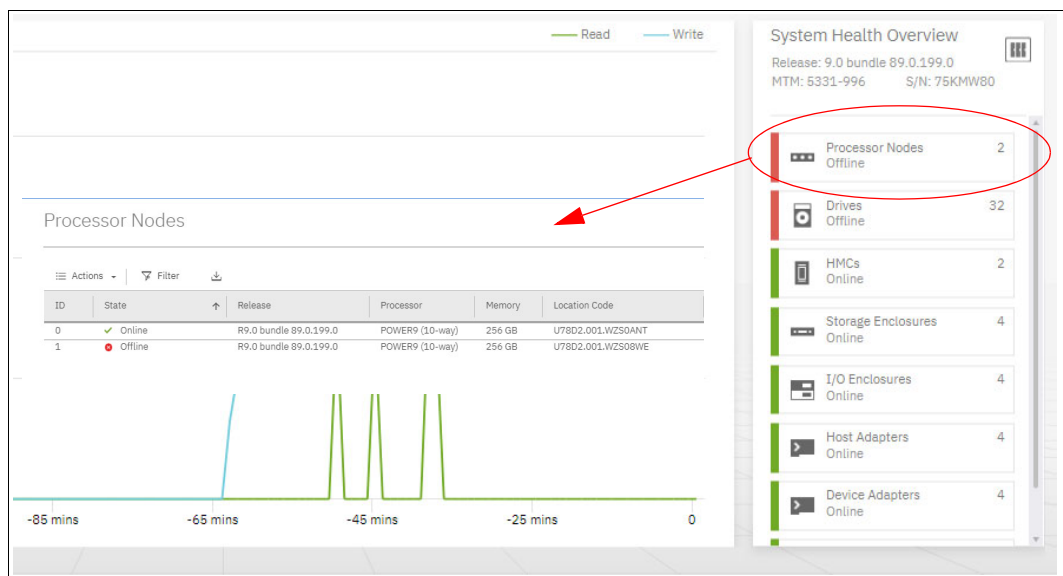


Figure 9-104 System Health Overview with a failed processor node

9.13.3 Monitoring system events

The Events window displays all events that occurred within the storage system, whether they are initiated by a user or by the system.

The Events table updates continuously so that you can monitor events in real time and track events historically.

Events are categorized by five levels of severity:

- ▶ Error
- ▶ Warning
- ▶ Inactive error
- ▶ Inactive warning
- ▶ Information

To access the Events window, click **Monitoring** → **Events**. The Events window can also be displayed by clicking the **Event Status** icon, as shown in Figure 9-105.




Figure 9-105 System Events window

The events can be exported as a CSV file by selecting **Export Table** on the Events window. The Export Table action creates a CSV file of the events that are displayed in the Events table with detailed descriptions.

9.13.4 Exporting system-wide information

As shown in Figure 9-106, you can export detailed reports from the dashboard with the following system-wide information into a CSV file.

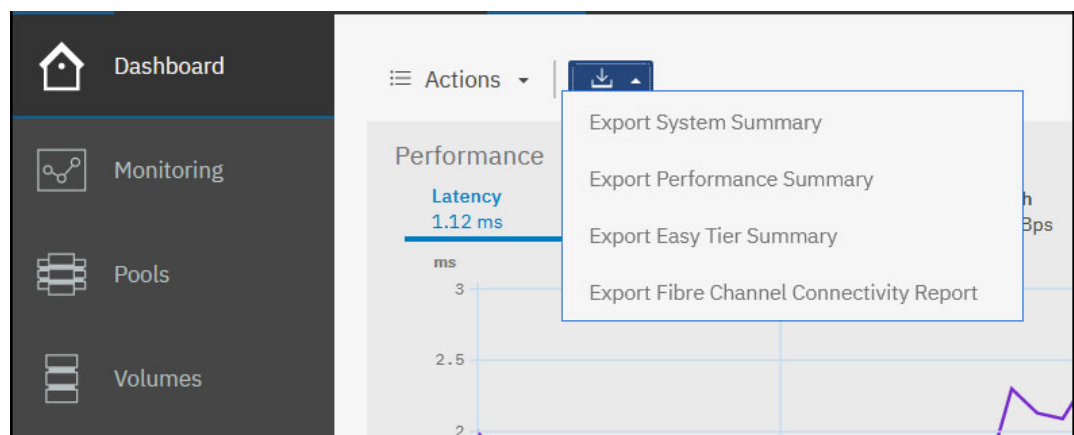


Figure 9-106 Export Performance Summary

► System Summary

This data includes CSV offload of the system summary (full hardware and logical configuration) by exporting individual tables (Volumes, I/O Ports, Users, and Volumes by LSS).

► Performance Summary

The performance data that is gathered at 1-minute intervals for up to one week can be exported for IBM Support analysis.

► Easy Tier Summary

You can directly offload the three CSV files and the Excel tool from both the DS GUI, as shown in Figure 9-111 on page 325, and the DS CLI (see 10.4.6, “IBM Easy Tier” on page 390). This download enables you to view the detailed data for Easy Tier planning, monitoring, and debugging.

Here are the offloaded CSV files:

- Workload categorization report: This report provides a method of classifying data activity that Easy Tier manages. This categorization includes a comparison of heat data across storage tiers.
- Workload skew curve report: This report provides a summary of workload distribution across tiers within a pool.
- Data movement report: This report provides details about Easy Tier data movement among the drive tiers on your storage system. Use the report to validate changes in tiers and the migration of warm and cool data.

For more information about Easy Tier, see *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667.

► Export Fibre Channel Connectivity Report

This report shows one row for every connected path between the DS8000 and a host system, a switch, or another DS8000. It also shows the status of these paths and their security settings. The report can be offloaded as a compressed file that contains a CSV file.

The CSV file name is IBM.2107-<serial>_fcconnectivity_<date>.csv, for example, IBM.2107-75DMC12_fcconnectivity_20210717.csv.

Example 9-1 shows a sample of the CSV file content, which can be imported and formatted in a spreadsheet tool such as Microsoft Excel.

Example 9-1 CSV content example

```
Local Port ID,LocalPort FC_ID,LocalPort WWPN,LocalPort WWNN,LocalPort Security
Capability,LocalPort Security Config,LocalPort Logins,LocalPort Security Capable
Logins,LocalPort Authentication Only Logins,LocalPort Encrypted Logins,AttachedPort
WWPN,AttachedPort WWNN,AttachedPort Interface ID,AttachedPort Type,AttachedPort
Model,AttachedPort Manufacturer,AttachedPort SN,RemotePort WWPN,RemotePort
WWNN,RemotePort FC_ID,RemotePort PRLI Complete,RemotePort Login Type,RemotePort Security
State,RemotePort Security Config,RemotePort Interface ID,RemotePort Type,RemotePort
Model,RemotePort SN,RemotePort SystemName
I0000,0x012E00,0x5005076306001339,0x5005076306FFD339,Auth Capable
Only,Disabled,1,0,0,0,0x202E8894715EC810,0x10008894715EC810,0x002E,008960,F64,IBM,000001
0550HA,0x5005076306041339,0x5005076306FFD339,0x010600,Yes,Mirroring secondary,
Notcapable,Disabled,0x0040,002107,996,0000000DMC01,Unknown
I0001,0x000002,0x5005076306005339,0x5005076306FFD339,Auth Capable
Only,Disabled,1,0,0,0,0x5005076306009339,0x5005076306FFD339,0x0002,002107,996,I
BM,0000000DMC01,0x5005076306009339,0x5005076306FFD339,0x000001,Yes,Mirroring
```

primary and
secondary, Notcapable, Disabled, 0x0002, 002107, 996, 0000000DMC01, Unknown
I0002, 0x000001, 0x5005076306009339, 0x5005076306FFD339, Auth Capable
Only, Disabled, 1, 0, 0, 0, 0x5005076306005339, 0x5005076306FFD339, 0x0001, 002107, 996, I
BM, 0000000DMC01, 0x5005076306005339, 0x5005076306FFD339, 0x000002, Yes, Mirroring
primary and
secondary, Notcapable, Disabled, 0x0001, 002107, 996, 0000000DMC01, Unknown
I0003, 0x000001, 0x500507630600D339, 0x5005076306FFD339, Auth Capable
Only, Disabled, 1, 0, 0, 0, 0x10000000C9CED91B, 0x20000000C9CED91B, 0x0000,

Figure 9-107 illustrates the formatting that occurs when you import the data in XLS format. Data is presented as one or more lines per port, depending on number of logins. This data illustrates four ports on one adapter, which are split into three captures for presentation.

Local port information									
Local Port ID	Local Port FC_ID	Local Port WWPN	Local Port WWNN	Local Port Security Capability	Local Port Security Config	Local Port Logins	Local Port Security Capable Logins	Local Port Authentication Only Logins	Local Port Encrypted Logins
I0000	0x012E00	0x5005076306001339	0x5005076306FFD339	Auth Capable Only	Disabled	1	0	0	0
I0001	0x000002	0x5005076306005339	0x5005076306FFD339	Auth Capable Only	Disabled	1	0	0	0
I0002	0x000001	0x5005076306009339	0x5005076306FFD339	Auth Capable Only	Disabled	1	0	0	0
I0003	0x000001	0x500507630600D339	0x5005076306FFD339	Auth Capable Only	Disabled	1	0	0	0

Attached port information						
Attached Port WWPN	Attached Port WWNN	Attached Port Interface ID	Attached Port Type	Attached Port Model	Attached Port Manufacturer	Attached Port SN
0x202E8894715EC810	0x10008894715EC810	0x002E	8960	F64	IBM	0000010550HA
0x5005076306009339	0x5005076306FFD339	0x0002	2107	996	IBM	0000000DMC01
0x5005076306005339	0x5005076306FFD339	0x0001	2107	996	IBM	0000000DMC01
0x10000000C9CED91B	0x20000000C9CED91B	Unknown	Unknown	Unknown	Unknown	Unknown

Remote port (Login) information											
Remote Port WWPN	Remote Port WWNN	Remote Port FC_ID	Remote Port PRLI Complete	Remote Port Login Type	Remote Port Security State	Remote Port Security Config	Remote Port Interface ID	Remote Port Type	Remote Port Model	Remote Port SN	Remote Port System Name
0x5005076306041339	0x5005076306FFD339	0x010600	Yes	Mirroring secondary	Not capable	Disabled	0x0040	2107	996	0000000DMC01	Unknown
0x5005076306009339	0x5005076306FFD339	0x000001	Yes	Mirroring primary and secondary	Not capable	Disabled	0x0002	2107	996	0000000DMC01	Unknown
0x5005076306005339	0x5005076306FFD339	0x000002	Yes	Mirroring primary and secondary	Not capable	Disabled	0x0001	2107	996	0000000DMC01	Unknown
0x10000000C9CED91B	0x20000000C9CED91B	0x000002	Yes	FCP host	Not capable	Disabled	Unknown	Unknown	Unknown	Unknown	Unknown

Figure 9-107 Example XLS format

9.13.5 Audit logs

Audit logs provide a record for auditing purposes to determine when changes were made to a storage system and by which user.

The audit log is an unalterable record of all actions and commands that were initiated by users on the system through the DS GUI, DS CLI, DS Network Interface (DSNI), or IBM Spectrum Control. The audit log does not include commands that were received from host systems or actions that were completed automatically by the storage system. The audit log is downloaded as a compressed text file.

An audit log for the DS8900F can also be exported by selecting **Export Audit Log** on the Events window, as shown in Figure 9-108.

Type	Time	Description
Warm start	10/04/2019 09:37:22 AM	Warm start occurred.
Warm start	10/04/2019 09:37:22 AM	Warm start occurred.
Warm start	10/02/2019 09:27:13 AM	Warm start occurred.
Warm start	10/02/2019 09:27:13 AM	Warm start occurred.
Warm start	10/01/2019 10:05:48 AM	Warm start occurred.
Warm start	10/01/2019 10:05:48 AM	Warm start occurred.
Warm start	10/01/2019 09:45:27 AM	Warm start occurred.
Warm start	10/01/2019 09:45:27 AM	Warm start occurred.
Warm start	10/01/2019 04:01:58 PM	Warm start occurred.
Warm start	10/01/2019 04:01:58 PM	Warm start occurred.
Warm start	09/27/2019 10:12:37 AM	Warm start occurred.
Warm start	09/27/2019 10:12:37 AM	Warm start occurred.
Warm start	09/26/2019 11:03:41 PM	Warm start occurred.
Warm start	09/26/2019 11:03:41 PM	Warm start occurred.
Warm start	09/23/2019 02:55:59 PM	Warm start occurred.
Warm start	09/23/2019 02:55:59 PM	Warm start occurred.
Warm start	09/21/2019 10:38:48 PM	Warm start occurred.
Warm start	09/21/2019 10:38:48 PM	Warm start occurred.

Figure 9-108 Export Audit Log

9.14 Performance monitoring

Performance data is presented directly on the initial dashboard view, or you can view it by clicking the **Monitoring** tab of the menu. To monitor the performance, complete the following steps:

1. From the System window, select the **Monitoring** icon.
2. Click **Performance** to display graphs that track performance indicators. The performance indicators include IOPS, latency, bandwidth, cache read hits, and write delays.
3. The Performance window includes four preset graphs, which are listed in the **Favorites** menu. These preconfigured graphs can be seen by clicking the **Favorites** menu icon, as shown in Figure 9-109. You also can set any of these graphs as the default.



Figure 9-109 Performance Monitoring window

You can create your own performance graphs for the storage system, pools, volumes, and FC ports. You can use predefined graphs and compare performance statistics for multiple pools, up to six volumes at a time, or FC ports.

Figure 9-110 shows a comprehensive view of the available performance functions.



Figure 9-110 Performance monitoring graph comprehensive view

To learn how to obtain statistics from the DS CLI, see 10.5, “Metrics with DS CLI” on page 391.

9.14.1 Performance statistics

Important: All the listed performance statistics are averaged over 1 minute. The performance graphs cover data that is collected for the last 7 days. For long-term performance statistics, use IBM Spectrum Control.

Users can monitor the following resources and performance metrics:

- ▶ System performance metrics:
 - IOPS: Total number of processed requests in thousand I/O operations per second (KIOPS) for the selected I/O operations.
 - Latency: Response time in milliseconds (ms) for the selected I/O operations.
 - Transfer Size: Number of kilobytes per I/O operation (KB/operation) for the selected I/O operations.
 - Bandwidth: Number of megabytes per second (MBps) for the selected bandwidth type.
 - Cache: Percentage of read hits (read I/O operations that were fulfilled from the system cache or memory) and write delays (I/O operations that were delayed because of write cache space constraints or other conditions) during 1 minute.
 - Capacity: Total capacity and provisioned capacity of the storage system in GiB.
 - Power: Power usage of the storage system in watts.
 - Temperature: Average temperature of the storage system in degrees Celsius.

- ▶ Array performance metrics:
 - Back-end IOPS: Number of requests that were processed on the selected arrays in KIOPS for the selected I/O operation type.
 - Back-end Latency: Response time in ms for the selected I/O operations that were processed on the chosen arrays.
 - Transfer Size: Number of KB/operation for the selected I/O operations (read, write, or average of both) that were processed on the selected arrays.
 - Back-end Bandwidth: Number of MBps for the selected bandwidth type (read, write, or combined total) that was processed on the selected arrays.
 - Utilization: Array utilization in %.
- ▶ Pool performance metrics:
 - Back-end IOPS: Number of requests that were processed on the arrays in the pool in KIOPS for the selected I/O operation type.
 - Back-end Latency: Response time in ms for the selected I/O operations that were processed on the arrays in the pool.
 - Transfer Size: Number of KB/operation for the selected I/O operations (read, write, or average of both) that were processed on the arrays in the pool.
 - Back-end Bandwidth: Number of MBps for the selected bandwidth type that was processed on the arrays in the selected pools.
 - Usable Capacity: Used, provisioned, and usable capacity, in GiB, in the pool.
 - Over-provisioned ratio: Over-provisioned capacity of the pool as a ratio of defined capacity over total capacity.
 - Data activity (Easy Tier): Capacity distribution of tiers in the pool showing capacity that is active, inactive, or in-between. You can click the bar graphs to zoom in and see a history of values.
 - Data movement (Easy Tier): Historical information showing the amount of data that is moved by Easy Tier within the pool.
 - Data policy (Easy Tier): The capacity of a pool's tiers that is pinned (assigned) to each tier.
- ▶ Volume performance metrics:
 - I/O: Measure the I/O in IOPS or Latency.
 - Bandwidth.
 - Data activity (Easy Tier): Capacity distribution of tiers in the volume showing capacity that is active, inactive, or in-between. You can click the bar graphs to zoom in and see a history of values.
 - Data policy (Easy Tier): The capacity of a volume's tiers that is pinned (assigned) to each tier.
- ▶ LSS performance metrics:
 - IOPS: Number of requests that are processed on the LSS in KIOPS for the default I/O operations (read, write, and total).
 - Latency: Response time in milliseconds (ms) for the default I/O operations (read, write, and average) that are processed on the LSS.
 - Bandwidth: Number of MBps for the selected bandwidth type (read, write, and total) that is processed on the LSS.

- ▶ HOST performance metrics:
 - IOPS: Number of requests that are processed on the host, in KIOPS for the default I/O operations (read, write, and total).
 - Latency: Response time in ms for the default I/O operations (read, write, and average) that is processed on the host.
 - Bandwidth: Number of MBps for the selected bandwidth type (read, write, and total) that is processed on the host.
- ▶ FC port performance metrics:
 - IOPS: Number of processed requests in KIOPS for the selected I/O operations (read, write, and total) on the FC port.
 - Latency: Response time in ms for the selected I/O operations on the FC port.
 - Transfer Size: Number of KB per operation for the selected I/O operations on the FC port.
 - Bandwidth: Number of MBps for the selected bandwidth type on the FC port.

You can use these performance metrics to define your own graphs. To add the custom graph to the Favorites menu, click the **star** icon, as shown in Figure 9-109 on page 322. You may also export the sample data that is used to create the performance graphs into a CSV file by clicking the **Save** icon, as shown in Figure 9-111.



Figure 9-111 StorageSystem_PerformanceSample export CSV file

9.14.2 Working with customized performance graphs

The DS GUI offers four predefined performance graphs, which are available in the Favorites menu, as shown in Figure 9-112.




Figure 9-112 Predefined graphs

For detailed performance analysis, you can define more detailed statistics and graphs, which can help identify and isolate problems. You can perform the following actions:

- ▶ Define your own performance graphs on demand.
- ▶ Add defined graphs to the Favorites menu.
- ▶ Pin defined graphs to the toolbar.
- ▶ Set defined graphs as a default in the Performance window.
- ▶ Rename or delete your graphs. You cannot delete predefined graphs.
- ▶ Change the time range of displayed graphs.

Creating array performance graphs

To create a graph from the System window, click the **Monitoring** icon. Click **Performance** to open the Performance Graphs window. Hover your cursor over the left bar and click to create a chart, as shown in Figure 9-113.




Figure 9-113 Creating a chart

Figure 9-114 demonstrates how to complete the following steps (each step number is referenced in the figure):

1. Select **Array** from the resources to monitor.
2. Select the arrays to monitor.
3. Select the metrics that you want (I/O, Bandwidth, or Utilization).

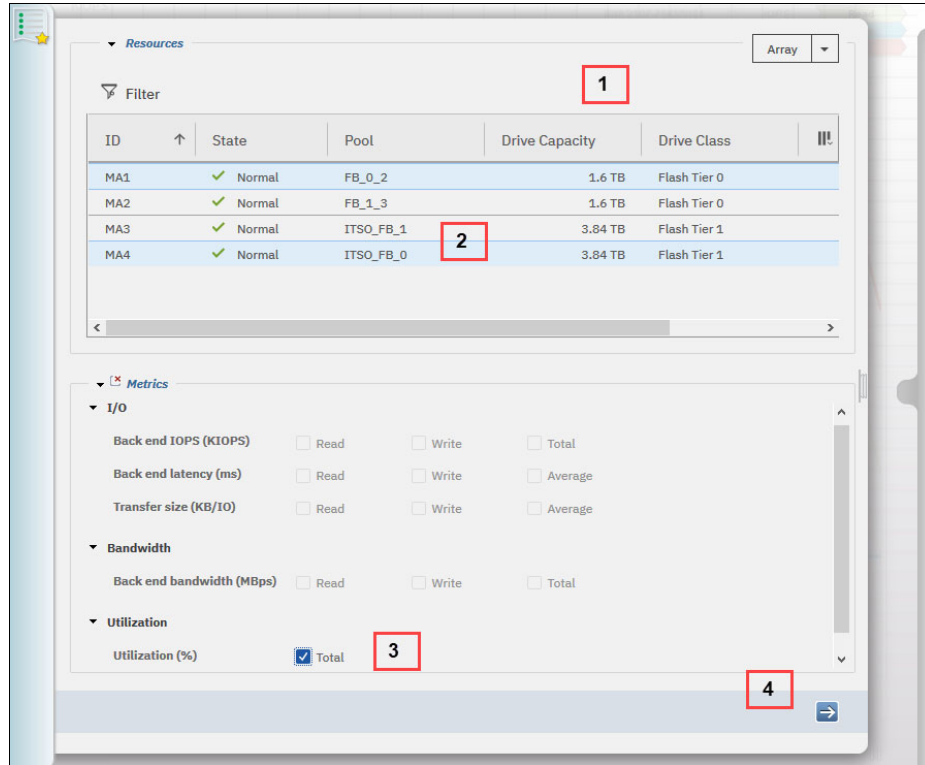


Figure 9-114 Array graph definition window

4. Draw the graph, which is shown in Figure 9-115.



Figure 9-115 Array performance graph: Total Utilization

Note: Array Utilization is the respective addition to Array performance metrics. When selected, it monitors the total utilization of the array.

To create a graph of a pool's performance, see Figure 9-113 on page 326, which shows how to create a chart, and then complete the following steps:

1. Select **Pool** from the resources to monitor.
2. Select the pool name to monitor.
3. Select the metrics that you want.

Figure 9-116 shows the metric options that are available for the selected pool.

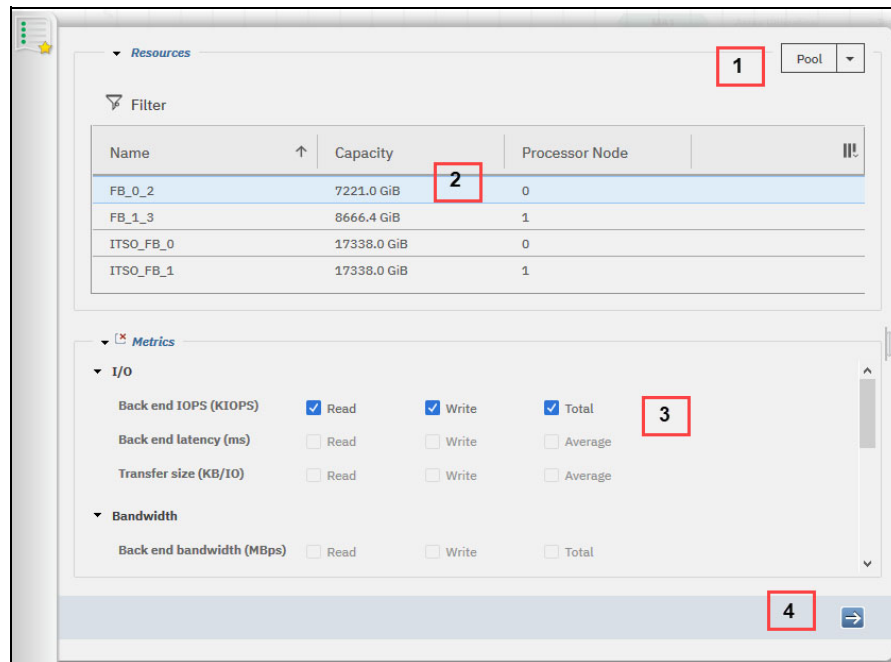


Figure 9-116 Pool graph definition window

4. Draw the graph, which is shown in Figure 9-117.




Figure 9-117 Pool performance graph

You can monitor Easy Tier directly in the DS GUI by using the workload categorization report and migration report. Figure 9-118 shows the Easy Tier pool level workload settings for creating total data movement report that is shown in Figure 9-119.

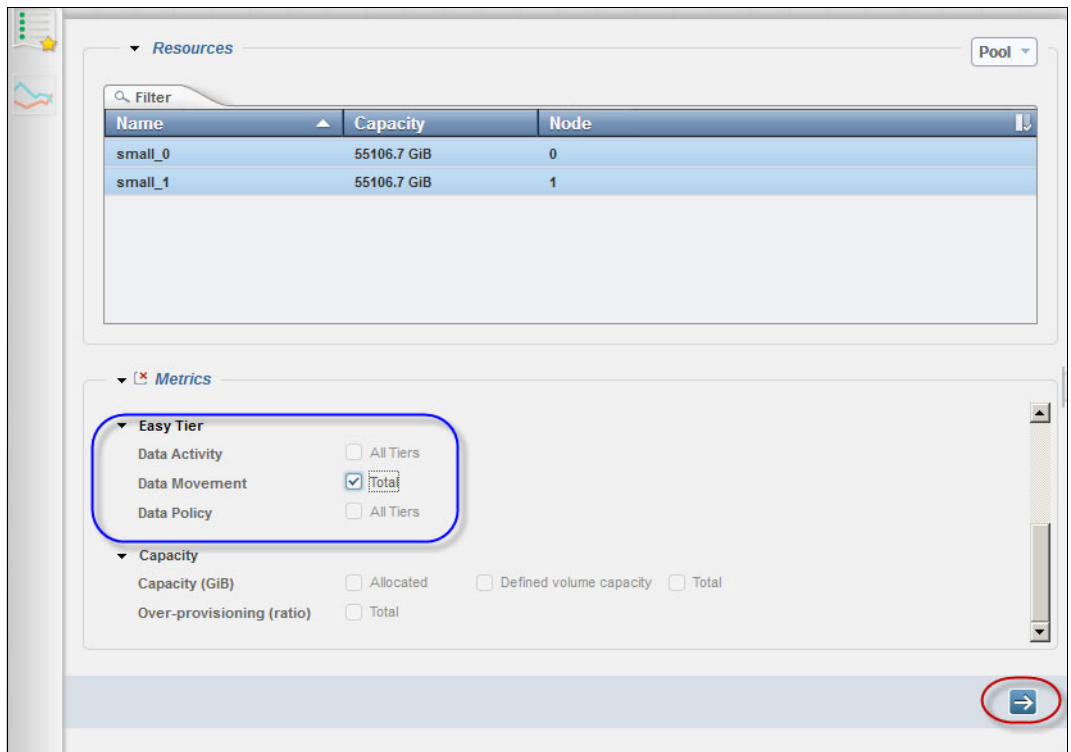


Figure 9-118 Example of Easy Tier settings for total Data Movement

Figure 9-119 shows an example of the Total Data Movement report.




Figure 9-119 Easy Tier pool level workload reports (Total Data Movement)

You can create an Easy Tier Data Activity Report for Pools and Volumes to monitor performance and to verify that data is placed in the right place. Select a pool to monitor, and then in the Easy Tier section select **All Tiers** for **Data Activity**, as shown in Figure 9-120.




Figure 9-120 Example of Easy Tier Data Activity Report setup

An example report of Easy Tier Data Activity for a pool is shown in Figure 9-121.




Figure 9-121 Example of Easy Tier Data Activity Report

Creating system performance graphs

Figure 9-113 on page 326 shows the steps to start a chart. To create the graph of the system's performance as shown in Figure 9-122, complete the following steps:

1. Select **System** from the resources to monitor.
2. Select the system to monitor.
3. Select the metrics that you want.




Figure 9-122 System graph definition window

4. Draw the graph, which is shown in Figure 9-123.




Figure 9-123 System performance graph

Creating FC port performance graphs

Start a chart as shown in Figure 9-113 on page 326. To complete the sequence that is indicated by the numbers in Figure 9-124 to create an FC port performance graph, complete the following steps:

1. Select **Fibre Channel Port** from the resources to monitor.
2. Select the system to monitor.
3. Select the metrics that you want.

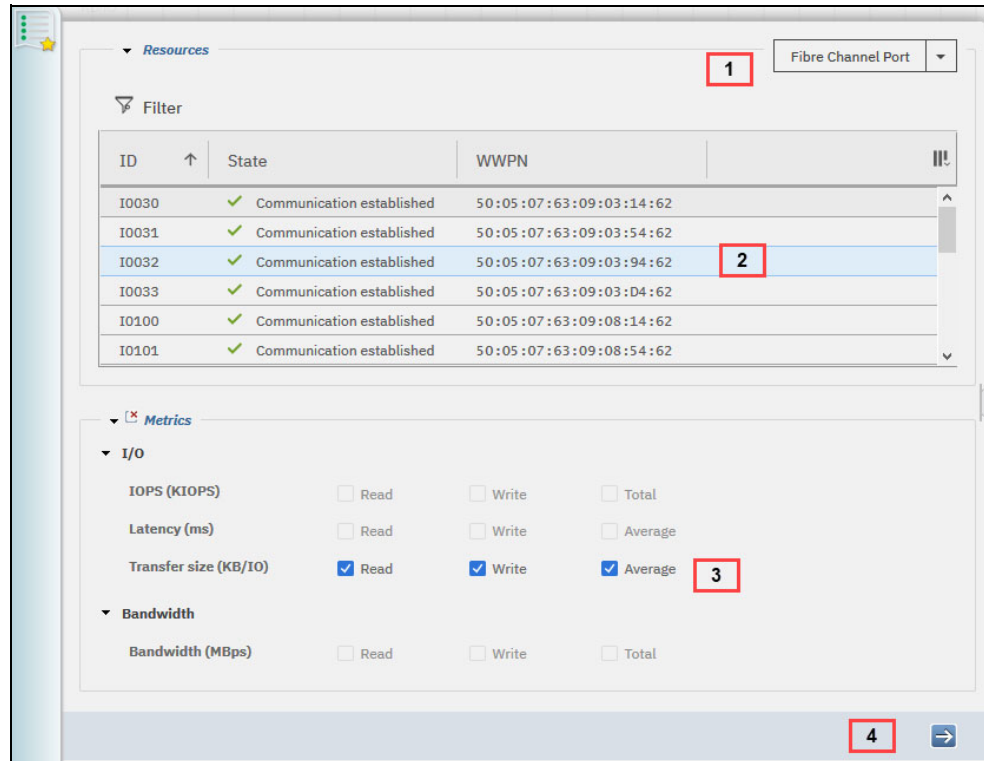


Figure 9-124 Fibre Channel Port graph definition window

4. Draw the graph. Figure 9-125 shows the completed graph.




Figure 9-125 Fibre Channel Port performance graph

Creating volume performance graphs

Start a chart as shown in Figure 9-113 on page 326. To create the Easy Tier performance graphs for the volume, as shown in Figure 9-126, complete the following steps:

1. Select **Volume** from the resources to monitor.
2. Select the volumes to monitor (you can select up to six volumes at a time for a graph).
3. Select the metrics that you want for I/O: **Bandwidth** or **Easy Tier**.

The screenshot shows the 'Volume Performance graph definition window' with the following components:

- Resources Section:** A table with columns 'Name', 'State', and 'Pool'. A dropdown menu labeled 'Volume' is at the top right, highlighted with a red box '1'. A 'Filter' icon is on the left. The table lists various storage volumes, with 'Team13_DB03' highlighted in blue, indicated by a red box '2'.
- Metrics Section:** A section titled 'Metrics' (highlighted with a red box '3') containing three categories of metrics:
 - I/O:** Includes 'IOPS (KIOPS)' (checked) and 'Latency (ms)' (unchecked).
 - Bandwidth:** Includes 'Bandwidth (MBps)' (unchecked).
 - Easy Tier:** Includes 'Data Activity' and 'Data Policy', both with 'All Tiers' (unchecked).
- Navigation:** A 'Next' arrow button is located at the bottom right, highlighted with a red box '4'.

Figure 9-126 Volume Performance graph definition window

4. Draw the graph. Figure 9-127 shows the bandwidth graph for a set of five volumes.




Figure 9-127 Volume performance graph

Note: The **Performance** action on the Volume, Host, and LSS resources is also available in all pages where they are shown. For any Volume, Host, or LSS, select **Performance** from the **Action** menu and then click the metric that you want to monitor. The performance window for the selected resource and metric opens. Figure 9-128 shows the performance actions and metrics that are available for a volume on the Volume by LSS page.

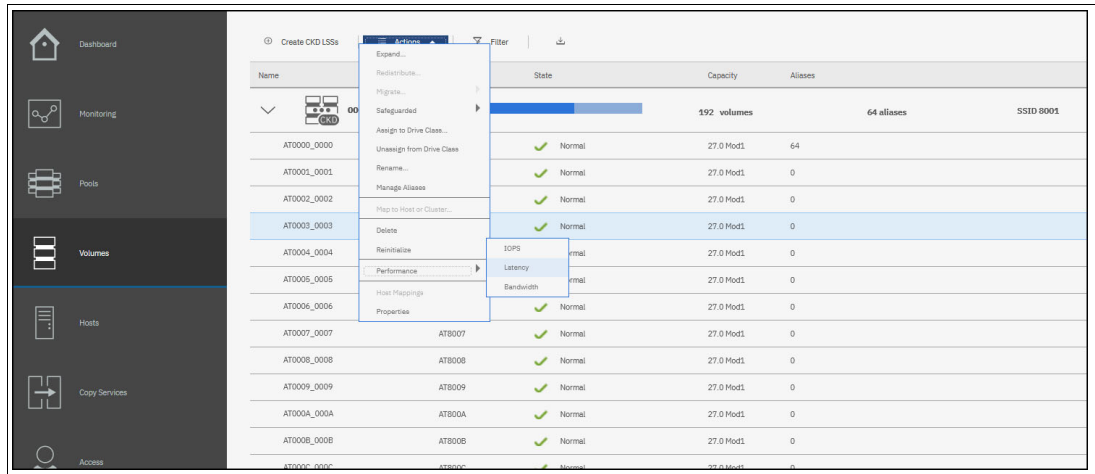


Figure 9-128 Performance action for a volume

Adding a graph to the Favorites menu

You can add a graph to the Favorites menu by completing the sequence that is indicated by the numbers that are shown in Figure 9-129 on page 335:

1. Click the icon **Add to Favorites**.
2. Enter a name for it.
3. Search and display your favorite graphs from the list.




Figure 9-129 Adding a graph to the Favorites menu

Pinning a graph to the toolbar

Graphs that are added to the Favorites menu can be pinned to the toolbar to enable faster access by completing the sequence that is indicated by the numbers that are shown in Figure 9-130.

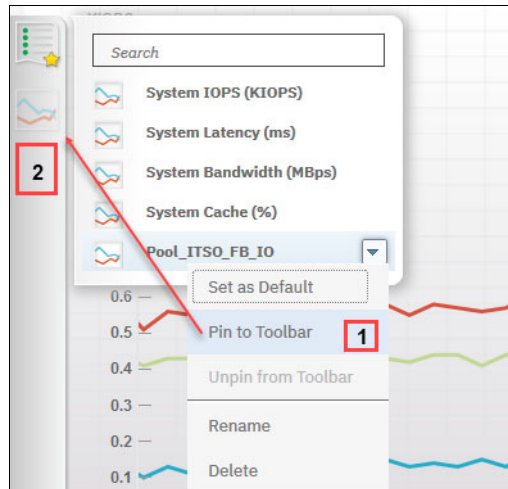


Figure 9-130 Pinning a graph to the toolbar

9.15 Fibre Channel error rate statistics

Errors that occurred during data transmission through FC ports over the past 24 hours can be displayed in the Error Rates for Fibre Channel Port window in the Network window, as shown in Figure 9-131. This feature is useful for a proactive check of SAN-related issues from a DS8900F perspective.

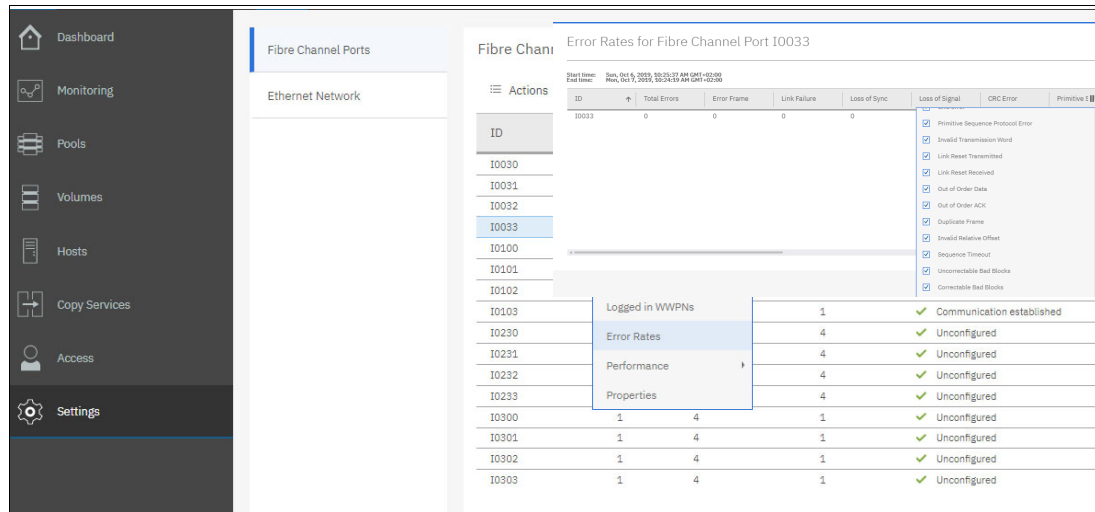


Figure 9-131 Fibre Channel Port error rate statistics

The following list shows all of the statistics that are available for port error checking:

- ▶ Total errors. The total number of errors that were detected on the FC port.
- ▶ Error frame: An FC frame was received that was not consistent with the FCP.
- ▶ Link failure: FC connectivity with the port was broken. This type of error can occur when the system that is connected to the port is restarted, replaced, or serviced, and the FC cable that is connected to the port is temporarily disconnected. It can also indicate a faulty connector or cable. Link failures result in degraded performance of the FC port until the failure is fixed.
- ▶ Loss of sync: A synchronization loss error was detected on the FC link. This type of error can occur when the system that is connected to the port is restarted, replaced, or serviced, and the FC cable that is connected to the port is temporarily disconnected. It also can indicate a faulty connector or cable. If a synchronization loss error persists, it can result in a link failure error.
- ▶ Loss of signal: A loss of signal was detected on the FC link. This type of error can occur when the system that is connected to the port is replaced or serviced, and the FC cable that is connected to the port is temporarily disconnected. It also can indicate a faulty connector or cable. If a loss of signal error persists, it can result in a link failure error.
- ▶ Cyclic redundancy check (CRC) error: An FC frame was received with CRC errors. This type of error is often fixed when the frame is retransmitted. This type of error is often recoverable and it does not degrade system performance unless the error persists and the data cannot be relayed after retransmission.
- ▶ Primitive sequence protocol error: A primitive sequence protocol error was detected. A primitive sequence is an ordered set that is transmitted and repeated continuously to indicate specific conditions within the port. The set also might indicate conditions that are encountered by the receiver logic of the port. This type of error occurs when an unexpected primitive sequence is received.

- ▶ **Transmission word count:** A bit error was detected. A transmission word is the smallest transmission unit that is defined in FC. This unit consists of four transmission characters, 4 x 10, or 40 bits. This type of error can include code violations, invalid special code alignment, and disparity errors.
- ▶ **Link reset transmitted:** The state of the FC port changed from active (AC) to link recovery (LR1).
- ▶ **Link reset received:** The state of the FC port changed from active (AC) to link recovery (LR2) state.
- ▶ **Out of order data:** A missing frame was detected. The frame was either missing from a data sequence or it was received beyond the FC port's sequence reassembly threshold.
- ▶ **Out of order acknowledgment (ACK):** An out of order ACK frame was detected. ACK frames signify that the transmission was received. The frame was either missing from a data sequence or it was received beyond the FC port's sequence reassembly threshold.
- ▶ **Duplicate frame:** A frame that was detected as previously processed was received.
- ▶ **Invalid relative offset:** A frame with an invalid relative offset parameter in the frame header was received.
- ▶ **Sequence timeout:** The FC port detected a timeout condition when a sequence initiator was received.
- ▶ **Uncorrectable bad blocks:** A data block with errors was unable to be fixed by Forward Error Correction (FEC).
- ▶ **Correctable bad blocks:** A data block with errors was fixed by FEC.
- ▶ **Transport mode write retries:** A transport mode write operation retry was requested. The buffer was not large enough to receive unsolicited data.

9.16 Providing feedback

Thirty days after the first login, a feedback request appears on the upper right under “Suggested tasks”, as shown in Figure 9-132. Click **Run Task** and a window opens that leads you through a survey. The task appears for all users in the system. After the feedback is complete, the system does not ask the particular user for further online feedback. After downloading the new bundle updates, it is not necessary to complete the survey again. If you click **Not now**, the feedback request can be snoozed for 90 days.




Figure 9-132 Suggested Tasks window



IBM DS8900F Storage Management Command-line Interface

This chapter describes how to use the DS8000 Command-line Interface (DS CLI).

This chapter covers the following topics:

- ▶ DS CLI overview
- ▶ I/O port configuration
- ▶ DS8900F storage configuration for Fixed-Block volumes
- ▶ DS8900F storage configuration for the CKD volumes
- ▶ Metrics with DS CLI
- ▶ Private network security commands
- ▶ Copy Services commands
- ▶ Earlier DS CLI commands and scripts
- ▶ For more information

Note: This chapter illustrates only a few essential commands. For a list of all commands and their parameters, see *IBM DS8000 Series Command-Line Interface User's Guide*, SC27-9562.

10.1 DS CLI overview

The DS CLI provides a full-function command set with which you can check your storage unit configuration and perform specific application functions. For more information about DS CLI usage and setup, see *IBM DS8000 Command-Line Interface User's Guide, SC27-9562*.

The following list highlights a few of the functions that you can perform with the DS CLI:

- ▶ Manage user IDs and passwords that can be used with DS GUI, DS CLI, and Hardware Management Console (HMC).
- ▶ Install activation keys for licensed features.
- ▶ Manage storage complexes and units.
- ▶ Configure and manage storage facility images (SFIs).
- ▶ Create and delete redundant array of independent disks (RAID) arrays, ranks, and extent pools.
- ▶ Create and delete logical volumes.
- ▶ Manage the host access to volumes.
- ▶ Check the current Copy Services (CS) configuration that is used by the storage unit.
- ▶ Create, modify, or delete CS configuration settings.
- ▶ Integrate Lightweight Directory Access Protocol (LDAP) policy usage and configuration.
- ▶ Implement encryption functions.

Single installation: In almost all cases, you can use a single installation of the current version of the DS CLI for all of your system needs. However, it is not possible to test every version of DS CLI with every Licensed Machine Code (LMC) level. Therefore, an occasional problem might occur despite every effort to maintain that level of compatibility.

If you suspect a version incompatibility problem, install the DS CLI version that corresponds to the LMC level that is installed on your system. You can have more than one version of DS CLI installed on your system, each in its own directory.

10.1.1 Supported operating systems for the DS CLI

The DS CLI can be installed on many operating systems (OSs):

- ▶ AIX
- ▶ Red Hat Linux
- ▶ SUSE Linux
- ▶ IBM i
- ▶ Oracle Solaris
- ▶ VMware ESX
- ▶ Microsoft Windows
- ▶ IBM z/OS

Important: For more information about supported OSs, specific preinstallation concerns, and installation file locations, see [IBM Documentation](#).

Before you can install the DS CLI, ensure that at least Java 8 or later is installed. A suitable level of Java might already be installed on many hosts. The installation program checks for this requirement during the installation process and does not install the DS CLI if a suitable version of Java is not already installed.

The installation process can be performed through a shell, such as the bash or Korn shell, the Windows command prompt, or through a GUI. If the installation process is installed by using a shell, the installation can be a silent installation by using a profile file. The installation process also installs software that allows the DS CLI to be uninstalled when the DS CLI is no longer required.

This chapter focuses on the DS CLI that is natively run from another system interacting with the DS8900F. For clarity purposes, the DS CLI is also available through the DS GUI by using the DS CLI option in the lower left of the dashboard.

10.1.2 Installation hints and tips

Before proceeding with the installation of the DS CLI, ensure that you are running the correct level of Java. If Java 8 or later is not found during the initial check, the following situation occurs:

- ▶ If you are using Windows, the following message is displayed:
LaunchAnywhere Error: Could not find a valid Java virtual machine to load.
You might need to reinstall a supported Java virtual machine.
- ▶ If you are using UNIX or Linux, the following message is displayed:
No Java virtual machine could be found from your PATH environment variable. You must install a VM before running this program.

After you ensure that Java 8 or later is installed, complete one of the following actions to correct the Java virtual machine Not Found error:

- ▶ Run the DS CLI installer again from the console, and provide the path to the Java virtual machine (JVM) by using the LAX_VM option. The following examples represent paths to the correct version of Java:
 - For a Windows system, specify the following path:
dsclicsetup.exe LAX_VM "C:\Program Files\java-whatever\jre\bin\java.exe"

Note: Due to a space in the Program Files directory name, you are required to add quotation marks around the directory name.

- For a UNIX or Linux system, specify the following path:
dsclicsetup.bin LAX_VM /opt/ibm-Java-whatever/java/bin/java

Note: If you use the LAX_VM argument, the installer attempts to use whatever JVM that you specify, even if it is an unsupported version. If an unsupported version is specified, the installation *might* complete successfully, but the DS CLI might not run and return an Unsupported Class Version Error message. You must ensure that you specify a supported version.

- Continue with the installation of the DS CLI.

- ▶ (For UNIX or Linux) Add the JVM location to your **PATH** environment variable by running the following command:

```
export PATH=$PATH:/opt/ibm-Java-whatever/java/bin
```

Then, run the **dsclicsetup.bin** program to install the DS CLI.

- ▶ (AIX only) Run the following commands to sequentially disable the **LIBPATH** environment variable, install the DS CLI, and restore the **LIBPATH** environment variable:

```
export LIBSAVE=$LIBPATH unset LIBPATH
dsclicsetup.bin LAX_VM/opt/ibm-Java-whatever/java/bin/java
export LIBPATH=$LIBSAVE
unset LIBSAVE
```

10.1.3 Installing the DS CLI on a Windows 10 or 11 system

To install the DS CLI on a Microsoft Windows 10 or 11 system, make sure that the correct Java version is installed. Example 10-1 shows how to install DS CLI on a Windows 10 system with the extra step of passing the JVM path location. You can download the DS CLI from [IBM Fix Central](#).

Example 10-1 DS CLI installation with a path to java.exe on a Windows 10 system

```
E:\IMAGES\HMC\Disk1\InstData\Windows\NoVM>dsclicsetup.exe LAX_VM "C:\Program Files (x86)\Java\jre1.8.0_161\bin\java.exe"
```

For instances where Java is already set up, the installation starts with running the **dsclicsetup.exe** program that is found in your installation media, as shown in Figure 10-1.


<input type="checkbox"/>	Name	Date modified	Type	Size
<input checked="" type="checkbox"/>	 dsclicsetup	2020-08-16 5:43 PM	Application	15,320 KB

Figure 10-1 The **dsclicsetup.exe** file

10.1.4 Installing the DS CLI on an z/OS system

You can install the DS CLI along with IBM Copy Services Manager on z/OS system. It is a regular System Modification Program/Extended (SMP/E) installation.

The DS CLI runs under UNIX System Services for z/OS, and has a separate FMID HIWN62M. You can also install the DS CLI separately from IBM Copy Services Manager.

For more information, see *IBM DS CLI on z/OS Program Directory*, G113-3563. You can use the order number (G113-3563) to search for it at [IBM Publications Center](#).

After the installation is done, the first thing to do is to access your UNIX System Services for z/OS. This process can vary from installation to installation. Ask your z/OS system programmer how to access it.

Tip: Set your Time Sharing Option (TSO) **REGION SIZE** to 512 MB to allow the DS CLI to run.

In our test system, we logged on to TSO and used option 6 ISPF Command Shell, we issued the command **OMVS** to start the z/OS UNIX Shell, as shown in Figure 10-2.

```
Menu List Mode Functions Utilities Help
-----
MCECEBC                                ISPF Command Shell
Enter TSO or Workstation commands below:

===> omvs

Place cursor on choice and press enter to Retrieve command

=> omvs
```

Figure 10-2 OMVS command to start the z/OS UNIX Shell

The default installation path for the z/OS DS CLI is /opt/IBM/CSMDCLI. To run the DS CLI, change your working directory to the installation path by issuing the following command, as shown in Figure 10-3:

```
cd /opt/IBM/CSMDCLI
```

```
IBM
Licensed Material - Property of IBM
...
GSA ADP Schedule Contract with IBM Corp.

IBM is a registered trademark of the IBM Corp.

-----
Business Notice:
  IBM's internal systems must only be used for conducting IBM's
  business or for purposes authorized by IBM management.
-----

$

===> cd /opt/IBM/CSMDCLI

INPUT
ESC=¢    1=Help    2=SubCmd    3=HlpRetrn  4=Top    5=Bottom    6=TSO    7=BackScr
8=Scroll  9=NextSess 10=Refresh 11=FwdRetr 12=Retrieve
```

Figure 10-3 The cd /opt/IBM/CSMDCLI command

You can choose between issuing the following commands to start DS CLI on z/OS, as shown in Figure 10-4:

- ▶ `./dscli`
- ▶ `./dscli -cfg ./profile/aca91.profile` (If you decide to use a profile.)

```
IBM
Licensed Material - Property of IBM
...
GSA ADP Schedule Contract with IBM Corp.

IBM is a registered trademark of the IBM Corp.
-----
IBM's internal systems must only be used for conducting IBM's
business or for purposes authorized by IBM management.
-----

$ cd /opt/IBM/CSMDSCLI
$
==> ./dscli

INPUT
ESC=¢    1=Help    2=SubCmd    3=HlpRetrn  4=Top        5=Bottom    6=TS0        7=BackScr
8=Scroll  9=NextSess 10=Refresh 11=FwdRetr  12=Retrieve
```

Figure 10-4 The `./dscli` command

If you change your mind and decide to quit here, instead of typing `./dscli`, press F2 to activate the SubCmd, as shown in Figure 10-4 (2=SubCmd). The OMVS Subcommand line is displayed, and you can issue a `quit` command.

As shown in Figure 10-5, the message CEE5210S The signal SIGHUP was received followed by `***` appears. Press Enter to quit OMVS.

```
IBM
Licensed Material - Property of IBM
...
IBM is a registered trademark of the IBM Corp.
...
-----

$ cd /opt/IBM/CSMDSCLI
$
OMVS Subcommand ==> quit

SUBCOMMAND
ESC=    1=Help    2=SubCmd 3=Return  4=Top        5=Bottom    6=TS0        7=BackScr
8=Scroll  9=NextSess 10=Refresh 11=FwdRetr  12=Retrieve

-----

CEE5210S The signal SIGHUP was received.
***
```

Figure 10-5 Sequence to leave the DS CLI

By using the DS CLI on z/OS, you can issue single commands, use a script mode, or go into batch mode by using z/OS Job Control Language (JCL). Figure 10-6 shows how to access the command interface.

```

Business Notice:
  IBM's internal systems must only be used for conducting IBM's
  business or for purposes authorized by IBM management.
-----

$ cd /opt/IBM/CSMDSCLI
$ ./dscli
Enter the primary management console IP address: <enter-your-machine-ip-address>
Enter the secondary management console IP address:
Enter your username: <enter-your-user-name-as-defined-on-the-machine>
Enter your password: <enter-your-user-password-to-access-the-machine>
dscli> ver -l
...
dscli>
===>

INPUT
ESC=¢    1=Help    2=SubCmd    3=HlpRetrn  4=Top        5=Bottom    6=TSO        7=BackScr   8=Scroll
9=NextSess 10=Refresh 11=FwdRetr  12=Retrieve

```

Figure 10-6 Accessing the DS CLI on z/OS

The command that you run on DS CLI on z/OS has the same syntax as in other platforms. Some examples of those commands are shown in Figure 10-7.

```

dscli> lssii
Name      ID          Storage Unit  Model WWNN          State ESSNet
=====
IBM.2107-75ACA91 IBM.2107-75ACA91 IBM.2107-75ACA90 980 5005076303FFD13E Online Enabled
dscli> lsckdvol -lcu EF
Name      ID  accstate  datastate  configstate  deviceMTM  voltype  orgbvols  extpool  cap (cyl)
=====
ITSO_EF00 EF00 Online    Normal    Normal    3390-A  CKD Base -  P1      262668
ITSO_EF01 EF01 Online    Normal    Normal    3390-9  CKD Base -  P1      10017
dscli> mkckdvol -dev IBM.2107-75ACA91 -cap 3339 -datatype 3390 -eam rotateexts -name ITSO_#h -extpool P1
EF02-EF02
CMUC00021I mkckdvol: CKD Volume EF02 successfully created.
dscli> lsckdvol -lcu EF
Name      ID  accstate  datastate  configstate  deviceMTM  voltype  orgbvols  extpool  cap (cyl)
=====
ITSO_EF00 EF00 Online    Normal    Normal    3390-A  CKD Base -  P1      262668
ITSO_EF01 EF01 Online    Normal    Normal    3390-9  CKD Base -  P1      10017
ITSO_EF02 EF02 Online    Normal    Normal    3390-3  CKD Base -  P1      3339
dscli> rmckdvol EF02
CMUC00023W rmckdvol: The alias volumes associated with a CKD base volume are automatically deleted
before deletion of the CKD base volume. Are you sure you want to delete CKD volume EF02? Y/n": y
CMUC00024I rmckdvol: CKD volume EF02 successfully deleted.
dscli>
===>

```

Figure 10-7 Common commands on the DS CLI

Here are some examples of how to take advantage of using your own JCL:

- ▶ Example 10-2 shows a job to run a DS CLI script (multiple commands that are stored in a UNIX file and in script mode).

Example 10-2 Running multiple commands from a UNIX file

```
//job--goes-here
/*****/
//*                                          */
//* Run a DS CLI script in z/OS              */
//*                                          */
//* DS CLI on z/OS = CSMDSCLI               */
//*                                          */
//* All output (stdin and stderr) is directed to the job log */
//* The actual UNIX command follows the SH statement, it can span */
//* more than one line                       */
//*                                          */
/*****/
//*
//PBXBAT EXEC PGM=BPXBATCH
//STDIN DD DUMMY
//STDOUT DD SYSOUT=*
//STDERR DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//STDENV DD DUMMY
//STDPARM DD *
SH
/opt/IBM/CSMDSCLI/dscli -cfg /opt/IBM/CSMDSCLI/profile/aca91.profile
                        -script /u/itso/dscli_script.txt
/*
```

- ▶ Example 10-3 shows a JCL to run several commands in a row, each in single-shot mode.

Example 10-3 Several single-shot mode commands

```
//job-card-goes-here
/*****/
//*                                          */
//* Run several DS CLI command in z/OS      */
//*                                          */
//* DS CLI on z/OS = CSMDSCLI               */
//*                                          */
//* All output (stdin and stderr) is directed to the job log */
//* The actual UNIX commands follow the SH statement, they can */
//* span more than one line                 */
//*                                          */
//* Commands are separated by semicolon     */
//*                                          */
/*****/
//*
//PBXBAT EXEC PGM=BPXBATCH
//STDIN DD DUMMY
//STDOUT DD SYSOUT=*
//STDERR DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//STDENV DD DUMMY
```

```
//STDPARM DD *
SH
echo "Command 1:";
/opt/IBM/CSMDSCSI/dscli -cfg /opt/IBM/CSMDSCSI/profile/aca91.profile
    ver -1;
echo "Command 2:";
/opt/IBM/CSMDSCSI/dscli -cfg /opt/IBM/CSMDSCSI/profile/aca91.profile
    lsrnk;
echo "Command 3:";
/opt/IBM/CSMDSCSI/dscli -cfg /opt/IBM/CSMDSCSI/profile/aca91.profile
    larray;
/*
```

10.1.5 DS CLI version

The **ver** command displays the version of the DS CLI client program, the HMC code level (*Storage Manager*), the HMC DS CLI version, the LMC version, and the code bundle version. The **ver** command uses the following parameters:

- s** (optional): The **-s** parameter displays the version of the DS CLI client program. You cannot use the **-s** and **-1** parameters together.
- 1** (optional): The **-1** parameter displays the versions of the DS CLI client program, Storage Manager, HMC code level, LMC, and the code bundle. You cannot use the **-1** and **-s** parameters together.
- cli** (optional): The **-cli** parameter displays the version of the DS CLI client program. Version numbers are in this format:
version.release.modification.fixlevel
- stmgr** (optional): The **-stmgr** parameter displays the version of the Storage Manager. This ID is not the Storage Manager GUI. This ID relates to the HMC code level information.
- lmc** (optional): The **-lmc** parameter displays the version of the LMC.

Example 10-4 shows an example of the **ver** command, where the customer uses a earlier DS CLI version.

Example 10-4 DS CLI version command

```
dscli>ver -1
DSCLI          7.9.30.136
HMC DSCLI     7.9.30.154
=====Version=====
Storage Image  LMC          Bundle Version
=====
IBM.2107-75HAL91 7.9.30.154 89.30.68.0
dscli>
```

The following DS CLI commands check the IBM Copy Services Manager version and install new software levels of IBM Copy Services Manager:

- ▶ The **lsoftware -l** command displays the IBM Copy Services Manager version that is installed on both HMCs, as shown in Example 10-5.

Example 10-5 Displaying the IBM Copy Services Manager version

```
dscli> lsoftware -l
Type Version                Status HMC
=====
CSM V6.3.2.1-a20220503-1401 Running 1
CSM V6.3.2.1-a20220503-1401 Running 2
```

- ▶ The **installsoftware** command is used to install a new version of IBM Copy Services Manager software on the HMC, as shown in Example 10-6.

Example 10-6 The installsoftware command output

```
dscli> installsoftware -loc
/home/hscroot/wht/csm-setup-6.3.2.1-linux-x86_64.bin -certloc
/home/hscroot/wht/csm-setup-6.3.2.1-linux-x86_64.bin.crt -type CSM -hmc 1
IBM DSCLI Version: 0.0.0.0 DS: IBM.2107-75DMC41
CMUC00294I installsoftware: Upload file successfully.
CMUC00294I installsoftware: Software CSM is successfully installed on HMC 1.
```

10.1.6 User accounts

DS CLI communicates with the DS8900F through the HMC. The primary or secondary HMC can be used. The DS CLI access is authenticated by using IBM Enterprise Storage Server Network Interface (IBM ESSNI), which is also referred to as the DS Network Interface (DSNI) on the HMC. The same user ID is used for DS CLI and DS GUI access. For more information about user accounts, see 6.5, “User management” on page 187.

The default user ID is `admin` and the password is `admin`. The system forces you to change the password at the first login. If you forget the `admin` password, a reset can be performed that resets the `admin` password to the default value.

10.1.7 User management by using the DS CLI

Although many administrative functions of the DS8900F are controlled by the storage administrator, you might want to define other users with different authorities for other limited functions.

The following commands are used to manage user IDs by using the DS CLI:

- ▶ **mkuser**

A user account that can be used with the DS CLI and the DS GUI is created by using this command. Example 10-7 shows the creation of a user that is called `JohnDoe`, which is in the `op_storage` group. The temporary password of the user is `passw0rd`. The user must use the `chpass` command when they log in for the first time.

Example 10-7 Using the mkuser command to create a user

```
dscli> mkuser -pw passw0rd -group op_storage JohnDoe
CMUC00133I mkuser: User JohnDoe successfully created.
```

► **rmuser**

An existing user ID is removed by using this command. Example 10-8 shows the removal of a user called JaneSmith.

Example 10-8 Removing a user

```
dscli> rmuser JaneSmith  
CMUC00135W rmuser: Are you sure you want to delete user JaneSmith? [y/n]:y  
CMUC00136I rmuser: User JaneSmith successfully deleted.
```

► **chuser**

Use this command to change the password or group (or both) of an existing user ID. It also can be used to unlock a user ID that was locked by exceeding the allowable login retry count. The administrator can also use this command to lock a user ID. In Example 10-9, we unlock the user, change the password, and change the group membership for a user that is called JohnDoe. The user must use the **chpass** command the next time that they log in.

Example 10-9 Changing a user by using the chuser command

```
dscli> chuser -unlock -pw time2change -group op_storage JohnDoe  
CMUC00134I chuser: User JohnDoe successfully modified.
```

► **lsuser**

By using this command, a list of all user IDs can be generated. Example 10-10 shows a list of three users, including the administrator account.

Example 10-10 Using the lsuser command to list users

```
dscli> lsuser  
Name          Group          State  
=====
```

JohnDoe	op_storage	active
secadmin	admin	active
admin	admin	active

► **showuser**

The account details of a user ID can be displayed by using this command. Example 10-11 lists the details of the user JohnDoe.

Example 10-11 Using the showuser command to list user information

```
dscli> showuser JohnDoe  
Name          JohnDoe  
Group         op_storage  
State         active  
FailedLogin   0  
DaysToExpire 365  
Scope         PUBLIC
```

► **managepwfile**

An encrypted password file that is placed onto the local machine is created or added by using this command. This file can be referred to in a DS CLI profile. You can run scripts without specifying a DS CLI user password in clear text. If you are manually starting the DS CLI, you can also refer to a password file with the **-pwfile** parameter. By default, the file is in the following directories:

- Windows: C:\Users*<User>*\dscli\security.dat
- Other than Windows: HOME/dscli/security.dat

Example 10-12 shows managing the password file by adding the user ID JohnDoe. The password is now saved in an encrypted file that is called security.dat.

Example 10-12 Using the managepwfile command

```
dscli> managepwfile -action add -name JohnDoe -pw passw0rd
CMUC00206I managepwfile: Record 10.0.0.1/JohnDoe successfully added to password
file C:\Users\Administrator\dscli\security.dat.
```

► **chpass**

By using this command, you can change two password policies: Password expiration (in days) and the number of failed logins that are allowed. Example 10-13 shows changing the expiration to 365 days and five failed login attempts.

Example 10-13 Changing the rules by using the chpass command

```
dscli> chpass -expire 365 -fail 5
CMUC00195I chpass: Security properties successfully set.
```

► **showpass**

The properties for passwords (Password Expiration days and Failed Logins Allowed) are listed by using this command. Example 10-14 shows that passwords are set to expire in 90 days and that four login attempts are allowed before a user ID is locked.

Example 10-14 Using the showpass command

```
dscli> showpass
Password Expiration  365 days
Failed Logins Allowed 5
Password Age         0 days
Minimum Length      6
Password History     4
```

10.1.8 DS CLI profile

To access the DS8900F with the DS CLI, you must provide certain information by using the **dscli** command. At a minimum, the IP address or hostname of the DS8900F HMC, a username, and a password are required. You can also provide other information, such as the output format for list commands, the number of rows for each page in the command-line output, and whether a banner is included with the command-line output.

If you create one or more profiles to contain your preferred settings, you do not need to specify this information every time that you use the DS CLI. When you start the DS CLI, you can specify a profile name by using the **dscli** command. You can override the values of the profile by specifying a different parameter value for the **dscli** command.

When you install the command-line interface software, a default profile is installed in the profile directory with the software. The file name is `dscli.profile`. For example, use `C:\Program Files (x86)\IBM\dscli` for Windows and `/opt/ibm/dscli/profile/dscli.profile` for AIX (UNIX) and Linux platforms.

The following options are available for using profile files:

- ▶ You can modify the system default profile `dscli.profile`.
- ▶ You can create a personal default profile by copying the system default profile as `<user_home>/dscli/profile/dscli.profile`. The default home directory `<user_home>` is in the following directories:
 - Windows system: `%USERPROFILE%`, which is usually `C:\Users\Administrator`
 - UNIX or Linux system: `$HOME`
- ▶ You can create specific profiles for different storage units and operations. Save the profile in the user profile directory, for example:
 - `%USERPROFILE%\IBM\DSCLI\profile\operation_name1`
 - `%USERPROFILE%\IBM\DSCLI\profile\operation_name2`

Default profile file: The default profile file that you created when you installed the DS CLI might be replaced every time that you install a new version of the DS CLI. It is a best practice to open the default profile and then save it as under a new file. You can then create multiple profiles and reference the relevant profile file by using the `-cfg` parameter.

The following example uses a different profile when it starts the DS CLI:

```
dscli -cfg newprofile.profile (or whatever name you gave to the new profile)
```

These profile files can be specified by using the DS CLI command parameter `-cfg <profile_name>`. If the profile name is not specified, the default profile of the user is used. If a profile of a user does not exist, the system default profile is used.

Two default profiles: If two default profiles are called `dscli.profile`, one profile in the default system's directory and one profile in your personal directory, your personal profile is loaded.

Profile change illustration

To edit the profile, complete the following steps. This sequence assumes that your `%userprofile%` is `C:\Users\Administrator`.

1. Use Windows Explorer to copy the profile folder from `C:\Program Files (x86)\IBM\dscli` to `C:\Users\Administrator\dscli`.
2. From the Windows desktop, double-click the **DS CLI** icon.
3. In the command window that opens, enter the following command:

```
cd C:\Users\Administrator\dscli\
```
4. In the profile directory, enter the **notepad dscli.profile** command, as shown in Example 10-15.

Example 10-15 Command prompt operation

```
C:\Users\Administrator\dscli>cd profile  
C:\Users\Administrator\dscli\profile>notepad dscli.profile
```

5. The notepad opens and includes the DS CLI profile. Add four lines. Examples of these lines are shown in bold in Example 10-16.

Default newline delimiter: The default newline delimiter is a UNIX delimiter, which can render text in the notepad as one long line. Use a text editor that correctly interprets UNIX line endings.

Example 10-16 DS CLI profile example

```
#
# DS CLI Profile
#
# Management Console/Node IP Addresses
# hmc1 and hmc2 are equivalent to -hmc1 and -hmc2 command options.
#hmc1:127.0.0.1
#hmc2:127.0.0.1

# Default target Storage Image ID
# "devid" and "remotedevid" are equivalent to
# "-dev storage_image_ID" and "-remotedev storage_image_ID" command
options, respectively.
#devid: IBM.2107-AZ12341
#remotedevid:IBM.2107-AZ12341

devid: IBM.2107-75HAL91
hmc1: 10.0.0.1
username: admin
pwfile: c:\mydir\75HAL91\pwfile.txt
```

Adding the serial number by using the **devid** parameter and adding the HMC IP address by using the **hmc1** parameter are suggested. These additions help you to avoid mistakes when you use more profiles, and you do not need to specify this parameter for certain **dsccli** commands that require it. Additionally, if you specify the **dsccli** profile for CS usage, the **remotedevid** parameter is suggested for the same reasons. To determine the ID of a storage system, use the **lssi** CLI command.

Add the username and an encrypted password file by using the **managepwfile** command. A password file that is generated by using the **managepwfile** command is placed in the `user_home_directory\dsccli/profile/security/security.dat` directory. Specify the location of the password file with the **pwfile** parameter.

Important: Be careful if you add multiple **devid** and HMC entries. Uncomment (remove the number sign (#) one entry at a time. If multiple **hmc1** or **devid** entries exist, the DS CLI uses the entry that is closest to the bottom of the profile.

The following customization parameters also affect **dsccli** output:

- **banner:** Date and time with the **dsccli** version are printed for each command.
- **header:** Column names are printed.
- **format:** The output format, which is specified as **default**, **xml**, **delim**, or **stanza**.
- **paging:** For interactive mode, this parameter breaks output after several rows (24, by default).

6. After you save your changes, use Windows Explorer to copy the updated profile from `C:\Users\Administrator\dsccli\profile` to `C:\Program Files (x86)\IBM\dsccli\profile`.

10.1.9 Configuring the DS CLI to use the second HMC

The second HMC can be specified on the CLI or in the profile file that is used by the DS CLI. To specify the second HMC in a command, use the `-hmc2` parameter, as shown in Example 10-17.

Example 10-17 Using the -hmc2 parameter

```
C:\Program Files (x86)\IBM\dsccli>dsccli -hmc1 10.0.0.1 -hmc2 10.0.0.5
Enter your user name: JohnDoe
Enter your password: xxxxx
DS: IBM.2107-75HAL91
dsccli>
```

Alternatively, you can modify the following lines in the `dsccli.profile` (or any profile) file:

```
# Management Console/Node IP Addresses
# hmc1 and hmc2 are equivalent to -hmc1 and -hmc2 command options.
hmc1:10.0.0.1
hmc2:10.0.0.5
```

After these changes are made and the profile is saved, the DS CLI automatically communicates through HMC2 if HMC1 becomes unreachable. By using this change, you can perform configuration and run CS commands with full redundancy.

Two HMCs: If you specify only one HMC in a DS CLI command (or profile), any changes that you make to users are still replicated onto the other HMC.

10.1.10 Command structure

All commands that are used by the CLI follow a basic structure of up to four components in each command's use. A DS CLI command consists of 1 - 4 types of components that are arranged in the following order:

1. The command name: Specifies the task that the DS CLI performs.
2. Flags: Flags are used to modify the command. They provide more information that directs the DS CLI to perform the command task in a specific way.
3. Flags parameter: Provides information that is required to implement the command modification that is specified by a flag.
4. Command parameters: Provide basic information that is necessary to perform the command task. When a command parameter is required, it is always the last component of the command, and is not preceded by a flag.

10.1.11 Using the DS CLI application

To issue commands to the DS8900F, you must first log in to the DS8900F through the DS CLI with one of the following command modes of execution:

- ▶ Single-shot command mode
- ▶ Interactive command mode
- ▶ Script command mode

Single-shot command mode

Use the DS CLI single-shot command mode if you want to issue an occasional command from the operating system (OS) shell prompt where you need special handling, such as redirecting the DS CLI output to a file. You also use this mode if you are embedding the command into an OS shell script.

You must supply the login information and the command that you want to process at the same time. To use the single-shot mode, complete the following steps:

1. At the OS shell prompt, enter one of the following commands:
 - `dsccli -hmc1 <hostname or ip address> -user <adm user> -passwd <pwd> <command>`
 - `dsccli -cfg <dsccli profile> -pwfile <security file> <command>`

Important: Avoid embedding the username and password into the profile. Instead, use the `-pwfile` command.

2. Wait for the command to process and display the results.

Example 10-18 shows the use of the single-shot command mode.

Example 10-18 Single-shot command mode

```
C:\Program Files (x86)\IBM\dsccli>dsccli -hmc1 10.10.10.1 -user admin -passwd <pwd> lsuser
Name          Group                State
=====
AlphaAdmin    admin                locked
AlphaOper     op_copy_services    active
BetaOper      op_copy_services    active
admin         admin                active
```

Important: When you are typing the command, you can use the hostname or the IP address of the HMC. When a command is run in single-shot mode, the user must be authenticated. The authentication process can take a considerable amount of time.

Interactive command mode

Use the DS CLI interactive command mode when you want to issue a few infrequent commands without needing to log on to the DS8900F for each command.

The interactive command mode provides a history function that simplifies repeating or checking earlier command usage.

To use the interactive command mode, complete the following steps:

1. Log on to the DS CLI application at the directory where it is installed.
2. Provide the information that is requested by the information prompts. The information prompts might not appear if you provided this information in your profile file. The command prompt switches to a `dsccli` command prompt.
3. Use the DS CLI commands and parameters. You are not required to begin each command with `dsccli` because this prefix is provided by the `dsccli` command prompt.
4. Use the `quit` or `exit` command to end interactive mode.

Interactive mode: In interactive mode for long outputs, the message `Press Enter To Continue` appears. The number of rows can be specified in the profile file. Optionally, you can turn off the paging feature in the profile file by using the `paging:off` parameter.

Example 10-19 shows using interactive command mode by using the profile DS8900F.profile.

Example 10-19 Interactive command mode

```
C:\Program Files (x86)\IBM\dsccli>dsccli -cfg DS8900F.profile

dsccli> lsarraysite -l
arsite DA Pair dkcap (10^9B) diskrpm State   Array diskclass  encrypt
=====
S1    8           1600.0  65000 Assigned A2   FlashTier0 supported
S2    8           1600.0  65000 Assigned A3   FlashTier0 supported
S3   10           3840.0  65000 Assigned A1   FlashTier1 supported
S4   10           3840.0  65000 Assigned A0   FlashTier1 supported
dsccli> lssi
Name                ID                Storage Unit      Model WWNN                State  ESSNet
=====
IBM.2107-75HAL91  IBM.2107-75HAL91  IBM.2107-75HAL90  996   5005076309FFD462  Online Enabled
dsccli>
```

Script command mode

Use the DS CLI script command mode if you want to use a sequence of DS CLI commands. If you want to run a script that contains only DS CLI commands, you can start the DS CLI in script mode. The script that DS CLI runs can contain only DS CLI commands.

Example 10-20 shows the contents of a DS CLI script file. The file contains only DS CLI commands, although comments can be placed in the file by using a number sign (#). Empty lines are also allowed. One advantage of using this method is that scripts that are written in this format can be used by the DS CLI on any OS on which you can install the DS CLI. Only one authentication process is needed to run all of the script commands.

Example 10-20 Example of a DS CLI script file

```
# Sample dsccli script file
# Comments can appear if hashed
lsarraysite -l
lsarray -l
lsrank -l
```

For script command mode, you can turn off the banner and header for easier output parsing. Also, you can specify an output format that might be easier to parse by your script.

Example 10-21 shows starting the DS CLI by using the **-script** parameter and specifying a profile and the name of the script that contains the commands from Example 10-20.

Example 10-21 Running a DS CLI file

```
C:\Program Files (x86)\IBM\dsccli>dsccli -cfg DS8900F.profile -script c:\ds8000.script
arsite DA Pair dkcap (10^9B) diskrpm State   Array diskclass  encrypt
=====
S1    8           1600.0  65000 Assigned A2   FlashTier0 supported
S2    8           1600.0  65000 Assigned A3   FlashTier0 supported
S3   10           3840.0  65000 Assigned A1   FlashTier1 supported
S4   10           3840.0  65000 Assigned A0   FlashTier1 supported
Array State      Data  RAIDtype  arsite Rank DA Pair DDMcap (10^9B) diskclass  encrypt
=====
A0  Assigned  Normal 6 (5+P+Q+S) S4    R0   10           3840.0 FlashTier1 supported
A1  Assigned  Normal 6 (5+P+Q+S) S3    R1   10           3840.0 FlashTier1 supported
A2  Assigned  Normal 6 (5+P+Q+S) S1    R2    8           1600.0 FlashTier0 supported
A3  Unassigned Normal 6 (5+P+Q+S) S2    -    8           1600.0 FlashTier0 supported
```

ID	Group	State	datastate	Array	RAIDtype	extpoolID	extpoolnam	stgtype	exts	usedexts	keygrp	marray	extsize (cap)
R0	0	Normal	Normal	A0	6	P0	ITS0_FB	fb	17338	2105	1	MA4	1GiB
R1	1	Normal	Normal	A1	6	P1	ITS0_FB	fb	17338	2304	1	MA3	1GiB
R2	-	Unassigned	Normal	A2	6	-	-	fb	7221	-	1	MA1	1GiB

Important: The DS CLI script can contain only DS CLI commands. Using shell commands results in a process failure.

10.1.12 Return codes

When the DS CLI exits, the exit status code is provided, which effectively is a return code. If DS CLI commands are issued as separate commands (rather than by using script mode), a return code is presented for every command. If a DS CLI command fails (for example, because of a syntax error or the use of an incorrect password), a failure reason and return code are shown. Standard techniques to collect and analyze return codes can be used.

The return codes that are used by the DS CLI are listed in *Command-Line Interface User's Guide*, SC27-9562.

10.1.13 User assistance

The DS CLI is designed to include several forms of user assistance. The main form of user assistance is through [IBM Documentation](#).

Click the **Command-line interface** tab to access user assistance. You can also get user assistance when using the DS CLI program by running the **help** command. The following examples of usage are included:

- help** Lists all the available DS CLI commands.
- help -s** Lists all the DS CLI commands with brief descriptions of each command.
- help -l** Lists all the DS CLI commands with their syntax information.

To obtain information about a specific DS CLI command, enter the command name as a parameter of the **help** command. The following examples of usage are included:

- help <command name>** Provides a detailed description of the specified command.
- help -s <command name>** Provides a brief description of the specified command.
- help -l <command name>** Provides syntax information about the specified command.

Man pages

A *man page* is available for every DS CLI command. Man pages are most commonly seen in UNIX OSs, and provide information about command capabilities. This information can be displayed by issuing the relevant command followed by the **-h**, **-help**, or **-?** flags.

10.2 I/O port configuration

Set the I/O ports to the topology that you want. Example 10-22 lists the I/O ports by using the `lsetioport` command. I0030 - I0033 are on one adapter, and I0100 - I0103 are on another adapter.

Example 10-22 Listing the I/O ports

```
dsccli> lsetioport
```

ID	WWPN	State	Type	topo	portgrp	Security
I0030	5005076309031462	Communication established	Fibre Channel-SW	SCSI-FCP	0	Disabled
I0031	5005076309035462	Communication established	Fibre Channel-SW	SCSI-FCP	0	Disabled
I0032	5005076309039462	Communication established	Fibre Channel-SW	SCSI-FCP	0	Disabled
I0033	500507630903D462	Communication established	Fibre Channel-SW	SCSI-FCP	0	Disabled
I0100	5005076309081462	Communication established	Fibre Channel-SW	SCSI-FCP	0	Disabled
I0101	5005076309085462	Communication established	Fibre Channel-SW	SCSI-FCP	0	Disabled
I0102	5005076309089462	Communication established	Fibre Channel-SW	SCSI-FCP	0	Disabled
I0103	500507630908D462	Communication established	Fibre Channel-SW	SCSI-FCP	0	Disabled
I0230	5005076309131462	No light detected	Fibre Channel-LW	FICON	0	Disabled
I0231	5005076309135462	No light detected	Fibre Channel-LW	FICON	0	Disabled
I0232	5005076309139462	No light detected	Fibre Channel-LW	FICON	0	Disabled
I0233	500507630913D462	No light detected	Fibre Channel-LW	FICON	0	Disabled
I0300	5005076309181462	Offline	Fibre Channel-LW	-	0	Disabled
I0301	5005076309185462	Offline	Fibre Channel-LW	-	0	Disabled
I0302	5005076309189462	Offline	Fibre Channel-LW	-	0	Disabled
I0303	500507630918D462	Offline	Fibre Channel-LW	-	0	Disabled

The following possible topologies for each I/O port are available:

- ▶ Small Computer System Interface - Fibre Channel Protocol (SCSI-FCP): Fibre Channel (FC)-switched fabric, which is also called *switched point-to-point*. This port type is also used for mirroring.
- ▶ Fibre Channel connection (IBM FICON): This port type is for IBM Z system hosts only.

Note: Fibre Channel Arbitrated Loop (FC-AL) is no longer an available topology because the host adapters in the DS8900F do not support it, but the switch **-topology fc-a1** is still there for compatibility with earlier versions.

The Security field indicates the status of the IBM Fibre Channel Endpoint Security feature. For more information, see *IBM Fibre Channel Endpoint Security for IBM DS8900F and IBM Z*, SG24-8455.

If added to the `setioport` command, the `-force` parameter allows a topology change to an online I/O port even if a topology is set. Example 10-23 shows setting I/O ports without and with the `-force` option to the FICON topology, and then checking the results.

Example 10-23 Changing the topology by using `setioport`

```
dsccli> lsetioport i0103
```

ID	WWPN	State	Type	topo	portgrp	Security
I0103	500507630908D462	Communication established	Fibre Channel-SW	SCSI-FCP	0	Disabled

```
dsccli> setioport -topology ficon I0103
CMUN04033E setioport: I0231: The change topology request failed because the current
topology is already set and force option is not provided.
dsccli> setioport -topology ficon -force I0103
CMUC00011I setioport: I/O Port I0231 successfully configured.
```

```

dscli> lsioport I0103
ID      WWPN              State              Type              topo  portgrp Security
=====
I0103  500507630908D462  Communication      established       Fibre Channel-SW FICON 0      Disabled

```

To monitor the status for each I/O port, see 10.5, “Metrics with DS CLI” on page 391.

10.3 DS8900F storage configuration for Fixed-Block volumes

This section reviews examples of a typical DS8900F storage configuration when the DS8900F storage is attached to open system hosts. You can perform the DS8900F storage configuration by completing the following steps:

1. Create the arrays.
2. Create the ranks.
3. Create the extent pools.
4. Optional: Create the repositories for space-efficient volumes (not included).
5. Create the volumes.
6. Create the volume groups.
7. Create the host connections.
8. Create a cluster and assign hosts to it.

10.3.1 Disk classes

Arrays and array sites are associated with the following disk classes (`diskclass`), which represent a classification of the different drive types that are available for the DS8900F:

- ▶ The flash Tier 0 flash disk class specifies 2.5-inch high-performance flash drives.
- ▶ The flash Tier 1 flash disk class specifies 2.5-inch high-capacity flash drives with a size of 3.84 TB.
- ▶ The flash Tier 2 flash disk class specifies 2.5-inch high-capacity flash drives with sizes of 1.92 TB, 7.68 TB, or 15.36 TB.

Important: For more information about the current drive choices and RAID capacities, see [IBM Documentation](#).

10.3.2 Creating the arrays

This step creates the arrays. Before the arrays are created, list the array sites. Use the `lsarraysite` command to list the array sites, as shown in Example 10-24 on page 359. *Array sites* are groups of eight drives that are predefined in the DS8900F.

Important: For a DS8900F, one *rank* is assigned to one *array*. An *array* is made of only one array site. An *array site* contains eight drives. There is a one to one relationship among array sites, arrays, and ranks.

Example 10-24 Listing array sites

```
dscli> lsarraysite -l
arsite DA Pair dkcap (10^9B) diskrpm State      Array diskclass  encrypt
=====
S1      8              1600.0   65000 Unassigned -      FlashTier0 supported
S2      8              1600.0   65000 Unassigned -      FlashTier0 supported
S3     10              3840.0   65000 Assigned  A1     FlashTier1 supported
S4     10              3840.0   65000 Assigned  A0     FlashTier1 supported
```

In Example 10-24, you can see two unassigned array sites. Therefore, you can create two arrays. The `-l` option reports the diskclass information.

You can issue the `mkarray` command to create arrays, as shown in Example 10-25. The example uses one array site to create a single RAID 6 array. If you want to create a RAID 10 array, change the `-raidtype` parameter to 10.

Example 10-25 Creating arrays by using mkarray

```
dscli> mkarray -raidtype 6 -arsite S1
CMUC00004I mkarray: Array A2 successfully created.
dscli> mkarray -raidtype 6 -arsite S2
CMUC00004I mkarray: Array A3 successfully created.
```

You can now see the arrays that were created by using the `lsarray` command, as shown in Example 10-26.

Example 10-26 Listing the arrays by using lsarray

```
dscli> lsarray
Array State      Data  RAIDtype      arsite Rank DA Pair DDMcap (10^9B)
=====
A0  Assigned  Normal 6 (5+P+Q+S) S4      R0  10              3840.0
A1  Assigned  Normal 6 (5+P+Q+S) S3      R1  10              3840.0
A2  Unassigned Normal 6 (5+P+Q+S) S1      -   8              1600.0
A3  Unassigned Normal 6 (5+P+Q+S) S2      -   8              1600.0
```

Example 10-27 shows the results of the `lsarraysite` command.

Example 10-27 Listing the high-performance flash disk class flash Tier 0 by using the lsarraysite command

```
dscli> lsarraysite -l -diskclass flashtier0
arsite DA Pair dkcap (10^9B) diskrpm State      Array diskclass  encrypt
=====
S1      8              1600.0   65000 Assigned A2     FlashTier0 supported
S2      8              1600.0   65000 Assigned A3     FlashTier0 supported
```

Example 10-28 shows the results of the `lsarray -l` command.

Example 10-28 Listing the disk class by using the `lsarray -l` command

```
dscli> lsarray -l
```

Array	State	Data	RAIDtype	arsite	Rank	DA	Pair	DDMcap (10 ⁹ B)	diskclass	encrypt
A0	Assigned	Normal	6 (5+P+Q+S)	S4	R0	10		3840.0	FlashTier1	supported
A1	Assigned	Normal	6 (5+P+Q+S)	S3	R1	10		3840.0	FlashTier1	supported
A2	Unassigned	Normal	6 (5+P+Q+S)	S1	-	8		1600.0	FlashTier0	supported
A3	Unassigned	Normal	6 (5+P+Q+S)	S2	-	8		1600.0	FlashTier0	supported

You can see the following information in the examples above:

- ▶ Type of RAID array (RAID 6).
- ▶ Number of drives that are allocated to the array (5+P+Q+S, which means that the usable space of the array is five times the drive size, and P+Q is used for R6 Parity and S for Spare).
- ▶ Capacity of the drives that are being used (3.84 TB and 1.6 TB).
- ▶ Array sites (S1 and S2) that were used to create the arrays.
- ▶ Disk class (FlashTier0 and FlashTier1).

Default RAID policy

RAID 6 is the recommended and default RAID type for all drives over 1 TB. You get an alert message if you try to use RAID 5 (RAID 5 is now supported only on the Flash Tier 0 disk class by using the Request for Price Quotation (RPQ) process), as shown in Example 10-29.

Example 10-29 Alert message about using RAID 5

```
dscli> mkarray -raidtype 5 -arsite S2 -force
CMUC00537I mkarray: You have accepted the following disclaimer: The use of RAID 6
over RAID 5 is highly recommended for increased reliability. Acknowledge that you
understand the risks associated with RAID 5.:
CMUN81160E mkarray: Cannot create the array because the capacity of the flash
drives prohibits use of RAID 5.
-----
with RPQ submitted and approved
-----
dscli> mkarray -raidtype 5 -arsite S2
CMUC00536W mkarray: The use of RAID 6 over RAID 5 is highly recommended for
increased reliability. Acknowledge that you understand the risks associated with
RAID 5.: Are you sure you want to accept the disclaimer above? [Y/N]: y
CMUC00004I mkarray: Array A3 successfully created.
dscli> lsarray
```

Array	State	Data	RAIDtype	arsite	Rank	DA	Pair	DDMcap (10 ⁹ B)
A0	Assigned	Normal	6 (5+P+Q+S)	S4	R0	10		3840.0
A1	Assigned	Normal	6 (5+P+Q+S)	S3	R1	10		3840.0
A2	Assigned	Normal	6 (5+P+Q+S)	S1	R2	8		1600.0
A3	Unassigned	Normal	5 (6+P+S)	S2	-	8		1600.0

10.3.3 Creating the ranks

After you create all of the required arrays, create the ranks by using the **mkrank** command. The format of the command is **mkrank -array Ax -stgtype xxx**, where **xxx** is Fixed-Block (FB) or Count Key Data (CKD), depending on whether you are configuring for open systems hosts or IBM Z system hosts.

After all the ranks are created, the **lsrank** command is run. This command displays this information:

- ▶ All the ranks that were created.
- ▶ The server to which the rank is attached (attached to none, in the example up to now).
- ▶ The RAID type.
- ▶ The format of the rank (fb or ckd).

Example 10-30 shows the **mkrank** command and the result of a successful **lsrank** command.

Example 10-30 Creating and listing ranks by using the mkrank and lsrank commands

```
dscli> mkrank -array A2 -stgtype fb
CMUC00007I mkrank: Rank R2 successfully created.
dscli> lsrank
ID Group State      datastate Array RAIDtype extpoolID stgtype
=====
R0  0 Normal    Normal    A0          6 P0        fb
R1  1 Normal    Normal    A1          6 P1        fb
R2  - Unassigned Normal    A2          6 -         fb
dscli> lsarray
Array State      Data  RAIDtype  arsite Rank DA Pair DDMcap (10^9B)
=====
A0  Assigned Normal 6 (5+P+Q+S) S4    R0  10          3840.0
A1  Assigned Normal 6 (5+P+Q+S) S3    R1  10          3840.0
A2  Assigned Normal 6 (5+P+Q+S) S1    R2   8          1600.0
A3  Unassigned Normal 6 (5+P+Q+S) S2    -   8          1600.0
dscli> mkrank -array A3 -stgtype fb -extsize 16mib
CMUC00007I mkrank: Rank R3 successfully created.
dscli> lsrank
ID Group State      datastate Array RAIDtype extpoolID stgtype
=====
R0  0 Normal    Normal    A0          6 P0        fb
R1  1 Normal    Normal    A1          6 P1        fb
R2  0 Normal    Normal    A2          6 P2        fb
R3  - Unassigned Normal    A3          5 -         fb
dscli> lsarray
Array State      Data  RAIDtype  arsite Rank DA Pair DDMcap (10^9B)
=====
A0  Assigned Normal 6 (5+P+Q+S) S4    R0  10          3840.0
A1  Assigned Normal 6 (5+P+Q+S) S3    R1  10          3840.0
A2  Assigned Normal 6 (5+P+Q+S) S1    R2   8          1600.0
A3  Assigned Normal 5 (6+P+S)  S2    R3   8          1600.0
```

When defining a rank, you can also specify the extent size. You can have ranks and extent pools with large 1 (gibabyte) GiB FB extents or small 16 mebibytes (MiB) FB extents. The extent unit is specified by the **-extsize** parameter of the **mkrank** command. The first rank that is added to an extent pool determines the extent size of the extent pool.

10.3.4 Creating the extent pools

The next step is to create the extent pools. Remember the following points when you create the extent pools:

- ▶ Each extent pool includes an associated rank group that is specified by the **-rankgrp** parameter, which defines the extent pool's server affinity (0 for server0 or 1 for server1).
- ▶ The extent pool type is FB or CKD, and is specified by the **-stgtype** parameter.
- ▶ The number of extent pools can range from one to the number of existing ranks. However, to associate ranks with both servers, you need at least two extent pools.

For easier management, create empty extent pools that relate to the type of storage or the planned usage for that pool. For example, create an extent pool pair for FB open systems environment and create an extent pool pair for the CKD environment.

When an extent pool is created, the system automatically assigns it an extent pool ID, which is a decimal number that starts from 0, preceded by the letter P. The ID that was assigned to an extent pool is shown in the CMUC00000I message, which is displayed in response to a successful **mkextpool** command.

Extent pools that are associated with rank group 0 receive an even ID number. Extent pools that are associated with rank group 1 receive an odd ID number. The extent pool ID is used when you refer to the extent pool in subsequent DS CLI commands. Therefore, it is best practice to note the ID.

Example 10-31 shows one example of extent pools that you can define on your system. This setup requires a system with at least four ranks.

Example 10-31 An extent pool layout plan

FB Extent Pool with 3840GB flash drives assigned to server 0 (FB_0)
FB Extent Pool with 3840GB flash drives assigned to server 1 (FB_1)
CKD Extent Pool with 1600GB flash drives assigned to server 0 (CKD_HPF_0)
CKD Extent Pool with 1600GB flash drives assigned to server 1 (CKD_HPF_1)

The **mkextpool** command forces you to name the extent pools. To do so, complete these steps:

1. Create empty extent pools by using the **mkextpool** command, as shown in Example 10-32.
2. List the extent pools to obtain their IDs.
3. Attach a rank to an empty extent pool by using the **chrank** command.
4. List the extent pools again by using **lsxtpool** and note the change in the capacity of the extent pool.

Example 10-32 Creating an extent pool by using mkextpool, lsxtpool, and chrank

```
dsccli> mkextpool -rankgrp 0 -stgtype fb FB_0
CMUC00000I mkextpool: Extent pool P2 successfully created.
dsccli> mkextpool -rankgrp 1 -stgtype fb FB_1
CMUC00000I mkextpool: Extent Pool P3 successfully created.
dsccli> lsxtpool
Name      ID stgtype rankgrp status availstor (2^30B) %allocated available reserved numvols
=====
FB_0     P2 fb          0 full          0          100          0          0          0
FB_1     P3 fb          1 full          0          100          0          0          0
dsccli> chrank -extpool P2 R2
```

CMUC00008I chrank: Rank R2 successfully modified.

dscli> **chrank -extpool P3 R3**

CMUC00008I chrank: Rank R3 successfully modified.

dscli> **lsextpool**

Name	ID	stgtype	rankgrp	status	availstor (2^30B)	%allocated	available	reserved	numvols
FB_0	P2	fb	0	below	7220	0	462099	64	0
FB_1	P3	fb	1	below	8665	0	554583	64	0

After a rank is assigned to an extent pool, you can see this change when you display the ranks.

In Example 10-33, you can see that rank R0 is assigned to extpool P0.

Example 10-33 Displaying the ranks after a rank is assigned to an extent pool

dscli> **lsrank**

ID	Group	State	datastate	Array	RAIDtype	extpoolID	stgtype
R0	0	Normal	Normal	A0	6 P0	fb	
R1	1	Normal	Normal	A1	6 P1	fb	
R2	0	Normal	Normal	A2	6 P2	fb	
R3	1	Normal	Normal	A3	5 P3	fb	
R8	0	Normal	Normal	A8	6 P4	ckd	
R11	1	Normal	Normal	A11	6 P5	ckd	

Example 10-34 shows the extent size in the query.

Example 10-34 Command showing the extent size

dscli> **lsrank -l**

ID	Group	State	datastate	Array	RAIDtype	extpoolID	extpoolnam	stgtype	exts	usedexts	keygrp	marray	extsize (cap)
R0	0	Normal	Normal	A0	6 P0	ITSO_FB	fb	17338	2105	1	MA4	1GiB	
R1	1	Normal	Normal	A1	6 P1	ITSO_FB	fb	17338	2304	1	MA3	1GiB	
R2	0	Normal	Normal	A2	6 P2	FB_0	fb	462163	0	1	MA1	16MiB	
R3	1	Normal	Normal	A3	5 P3	FB_1	fb	554647	0	1	MA2	16MiB	
R8	0	Normal	Normal	A8	6 P4	CKD_L	ckd	2392	31	1	MA25	1113cy1	
R11	1	Normal	Normal	A11	6 P5	CKD_L	ckd	8313	814	1	MA6	1113cy1	

10.3.5 Creating the FB volumes

Now, you can create volumes and volume groups. When you create the volumes or groups, try to distribute them evenly across the two rank groups in the storage unit.

Although an FB-type volume can be created as standard (thick) and thin (extent space efficient (ESE)-type volumes, this section describes the creation of the standard type only.

Creating the standard volumes

Use the following command format when creating a volume:

```
mkfbvol -extpool pX -cap xx -name high_fb_0#h XXXX-XXXX
```

The last parameter is the *volume_ID*, which can be a range or single entry. The four-digit entry is based on *LL* and *VV*. *LL* (00 - FE) equals the logical subsystem (LSS) that the volume belongs to, and *VV* (00 - FF) equals the volume number on the LSS. Therefore, the DS8900F can support 255 LSSs, and each LSS can support a maximum of 256 volumes.

Example 10-35 shows the creation of eight volumes, each with a capacity of 10 GiB. The first four volumes are assigned to rank group 0, and are assigned to LSS 20 with volume numbers 00 - 03. The second four volumes are assigned to rank group 1, and are assigned to LSS 21 with volume numbers of 00 - 03.

Example 10-35 Create fixed-block volumes by using mkfbvol

```

dscli> lsxtpool
Name      ID stgtype rankgrp status availstor (2^30B) %allocated available reserved numvols
=====
FB_0     P2 fb           0 below      7220         0  462099      64     0
FB_1     P3 fb           1 below      8665         0  554583      64     0
dscli> mkfbvol -extpool p2 -cap 10 -name fb_0_#h 2000-2003
CMUC00025I mkfbvol: FB volume 2000 successfully created.
CMUC00025I mkfbvol: FB volume 2001 successfully created.
CMUC00025I mkfbvol: FB volume 2002 successfully created.
CMUC00025I mkfbvol: FB volume 2003 successfully created.
dscli> mkfbvol -extpool p3 -cap 10 -name fb_1_#h 2100-2103
CMUC00025I mkfbvol: FB volume 2100 successfully created.
CMUC00025I mkfbvol: FB volume 2101 successfully created.
CMUC00025I mkfbvol: FB volume 2102 successfully created.
CMUC00025I mkfbvol: FB volume 2103 successfully created.

```

Looking closely at the **mkfbvol** command that is used in Example 10-35, you see that volumes 2000 - 2003 are in extpool P2. That extent pool is attached to rank group 0, which means server 0. Rank group 0 can contain only even-numbered LSSs, which means that volumes in that extent pool must belong to an even-numbered LSS. The first two digits of the volume serial number are the LSS number. So, in this case, volumes 2000 - 2003 are in LSS 20.

For volumes 2100 - 2103 in extpool P3 in Example 10-35, the first two digits of the volume serial number are 21 (an odd number), which signifies that they belong to rank group 1. The **-cap** parameter determines the size. However, because the **-type** parameter was not used, the default type is GiB or ds, which is a binary size of 2³⁰ bytes.

Therefore, these volumes are 10 GiB binary, which equates to 10,737,418,240 bytes. If you used the **-type ess** parameter, the volumes are decimally sized, and they are a minimum of 10,000,000,000 bytes in size.

Example 10-35 named the volumes by using the naming scheme fb_0_#h, where #h means that you are using the hexadecimal volume number as part of the volume name. This naming convention is shown in Example 10-36, where you list the volumes that you created by using the **lsfbvol** command. You then list the extent pools to see how much space is left after the volume is created.

Example 10-36 Checking the machine after the volumes are created by using lsfbvol and lsxtpool

```

dscli> lsfbvol
Name      ID  accstate  datastate  configstate  deviceMTM  datatype  extpool  cap (2^30B)  cap (10^9B)  cap (blocks)
=====
fb_0_2000 2000 Online    Normal     Normal      2107-900  FB 512   P2         10.0         -            20971520
fb_0_2001 2001 Online    Normal     Normal      2107-900  FB 512   P2         10.0         -            20971520
fb_0_2002 2002 Online    Normal     Normal      2107-900  FB 512   P2         10.0         -            20971520
fb_0_2003 2003 Online    Normal     Normal      2107-900  FB 512   P2         10.0         -            20971520

```



```
fb_1_2100  2100 Online Normal Normal 2107-900 FB 512 P3 10.0 - 20971520
fb_1_2101  2101 Online Normal Normal 2107-900 FB 512 P3 10.0 - 20971520
fb_1_2102  2102 Online Normal Normal 2107-900 FB 512 P3 10.0 - 20971520
fb_1_2103  2103 Online Normal Normal 2107-900 FB 512 P3 10.0 - 20971520
```

```
dscli> lsextpool
```

```
Name      ID stgtype rankgrp status availstor (2^30B) %allocated available reserved numvols
-----
FB_0     P2 fb          0 below      7180         0 459531      64    4
FB_1     P3 fb          1 below      8625         0 552015      64    4
```

Important considerations:

- ▶ For a DS8000, the LSSs can be ID 00 - FE. The LSSs are in address groups. Address group 0 is LSSs 00 - 0F, address group 1 is LSSs 10 - 1F, and so on, except group F, which is F0 - FE. When you create an FB volume in an address group, that entire address group can be used only for FB volumes. Be aware of this fact when you plan your volume layout in a mixed FB and CKD DS8000. The LSS is automatically created when the first volume is assigned to it.
- ▶ The `-perfgrp <perf_group_ID>` flag option is still available on the create volume commands for compatibility with earlier version, but the feature for Performance I/O Priority Manager was discontinued as of Release 9 DS8000 products.

Resource group: You can configure a volume to belong to a certain resource group by using the `-resgrp <RG_ID>` flag in the `mkfbvol` command. For more information, see *IBM System Storage DS8000 Copy Services Scope Management and Resource Groups*, REDP-4758.

T10 Data Integrity Field volumes

A standard for end-to-end error checking from the application to the disk drives is emerging that is called *SCSI T10 Data Integrity Field (DIF)*. T10 DIF requires volumes to be formatted in 520-byte sectors with cyclic redundancy check (CRC) bytes that are added to the data. If you want to use this technique, you must create volumes that are formatted for T10 DIF usage.

You configure T10 DIF by adding the `-t10dif` parameter to the `mkfbvol` command. It is possible to create T10 DIF volumes and use them as standard volumes, and then enable them later without configuration changes.

Storage pool striping

When a volume is created, you can choose how the volume is allocated in an extent pool with several ranks. The extents of a volume can be kept together in one rank (if enough free space exists on that rank). The next rank is used when the next volume is created. This allocation method is called *rotate volumes*.

You can also specify that you want the extents of the volume that you create to be evenly distributed across all ranks within the extent pool. This allocation method is called *rotate extents*. The storage pool striping spreads the I/O of a logical unit number (LUN) to multiple ranks, which improves performance and greatly reduces hot spots.

The extent allocation method (EAM) is specified by the **-eam rotateexts** or **-eam rotatevols** option of the **mkfbvol** command, as shown in Example 10-37.

Default allocation policy: For DS8900F, the default allocation policy is rotate extents.

Example 10-37 Creating a volume with storage pool striping

```
dscli> mkfbvol -extpool p2 -cap 15 -name fb_0_2004 -eam rotateexts 2004
CMUC00025I mkfbvol: FB volume 2004 successfully created.
```

The **showfbvol** command with the **-rank** option (Example 10-38) shows that the volume that you created is distributed across two ranks. It also shows how many extents on each rank were allocated for this volume. Compared to the examples above, the extent pool P2 now consists of two ranks, R2 and R3.

Example 10-38 Using showfbvol for information about a striped volume

```
dscli> showfbvol -rank 2004
Name                fb_0_2004
ID                  2004
acstate             Online
datastate           Normal
configstate         Normal
deviceMTM           2107-900
datatype            FB 512
addrgrp             2
extpool             P2
exts                960
cap (MiB)           15360
captype             DS
cap (2^30B)         15.0
cap (10^9B)         -
cap (blocks)        31457280
volgrp              -
ranks              2
dbexts              0
sam                 Standard
repcapalloc         -
eam                rotateexts
reqcap (blocks)     31457280
realextents         960
virtualextents     3
realcap (MiB)       15360
migrating           0
migratingcap (MiB) 0
perfgrp            PGO
migratingfrom       -
resgrp             RGO
tierassignstatus    Unknown
tierassignerror     -
tierassignorder     Unknown
tierassigntarget    Unknown
%tierassigned       0
etmonpauseremain   -
etmonitorreset      unknown
GUID                6005076309FFD46200000000000002004
safeguardedcap (2^30B) -
safeguardedloc     -
usedsafeguardedcap (2^30B) -
```

```

safeguarded                no
SGC Recovered              no
=====Rank extents=====
rank extents capacity (MiB/cyl) metadata
=====
R2      480 7680              yes
R3      480 7680              no

```

Dynamic Volume Expansion

A volume can be expanded without removing the data within the volume. You can specify a new capacity by using the **chfbvol** command, as shown in Example 10-39.

Example 10-39 Expanding a striped volume

```

dscli> chfbvol -cap 60 2004
CMUC00332W chfbvol: Some host operating systems do not support changing the volume
size. Are you sure that you want to resize the volume? [y/n]: y
CMUC00026I chfbvol: FB volume 2004 successfully modified.

```

The largest LUN size is 16 TiB. CS is not supported for LUN sizes larger than 4 TiB.

New capacity: The new capacity must be larger than the previous capacity. You cannot shrink the volume.

Because the original volume included the **rotateexts** attribute, the other extents are also striped, as shown in Example 10-40. See both examples and check the difference.

Example 10-40 Checking the status of an expanded volume

```

ddscli> showfbvol -rank 2004
Name                fb_0_2004
ID                  2004
accstate            Online
datastate           Normal
configstate         Normal
deviceMTM           2107-900
datatype            FB 512
addrgrp             2
extpool             P2
exts                3840
cap (MiB)           61440
captype             DS
cap (2^30B)         60.0
cap (10^9B)         -
cap (blocks)        125829120
volgrp              -
ranks              2
dbexts              0
sam                 Standard
repcapalloc         -
eam                 rotateexts
reqcap (blocks)     125829120
realextents         3840
virtualextents     7
realcap (MiB)       61440

```

```

migrating                0
migratingcap (MiB)       0
perfgrp                  PG0
migratingfrom            -
resgrp                   RGO
tierassignstatus         Unknown
tierassignerror          -
tierassignorder          Unknown
tierassigntarget         Unknown
%tierassigned            0
etmonpauseremain        -
etmonitorreset           unknown
GUID                    6005076309FFD4620000000000002004
safeguardedcap (2^30B)  -
safeguardedloc           -
usedsafeguardedcap (2^30B) -
safeguarded              no
SGC Recovered            no
=====Rank extents=====
rank extents capacity (MiB/cyl) metadata
=====
R2      1920 30720          yes
R3      1920 30720          yes

```

Important: Before you can expand a volume, you must delete all CS relationships for that volume.

Defining thin-provisioned ESE volumes

The DS8900F supports thin-provisioned volumes. A thin-provisioned volume is created by the DS CLI specifying the storage allocation method (SAM) of the ESE for a volume. This specification is done with the **-sam ese** option of the **mkfbvol** command, as shown in Example 10-41.

Example 10-41 Creating a thin-provisioned ESE volume

```

dscli> mkfbvol -extpool p2 -cap 1000 -name fb_0_2005 -sam ese 2005
CMUC00025I mkfbvol: FB volume 2005 successfully created.

```

The extent size is defined by the extent pools where the volume is created.

More DS CLI commands are available to control and protect the space in an extent pool for thin-provisioned volumes. One of these commands is the **mksestg** command, which reserves space for thin-provisioned volumes. For more information about thin-provisioning, see *IBM DS8880 Thin Provisioning (Updated for Release 8.5)*, REDP-5343.

Deleting volumes

FB volumes can be deleted by using the **rmfbvol** command. The command includes options to prevent the accidental deletion of volumes that are in use. An FB volume is considered to be in use if it is participating in a CS relationship or if the volume received any I/O operation in the previous 5 minutes.

Volume deletion is controlled by the **-safe** and **-force** parameters (they cannot be specified at the same time) in the following manner:

- ▶ If none of the parameters are specified, the system performs checks to see whether the specified volumes are in use. Volumes that are not in use are deleted and the volumes that are in use are not deleted.
- ▶ If the **-safe** parameter is specified and if any of the specified volumes are assigned to a user-defined volume group, the command fails without deleting any volumes.
- ▶ The **-force** parameter deletes the specified volumes without checking whether they are in use.

Example 10-42 shows the creation of volumes 2200 and 2201, and then the assignment of volume 2200 to a volume group. You try to delete both volumes with the **-safe** option, but the attempt fails without deleting either of the volumes. You can delete volume 2201 by using the **-safe** option because the volume is not assigned to a volume group. Volume 2200 is not in use, so you can delete it by not specifying either parameter.

Example 10-42 Deleting an FB volume

```
dsccli> mkfbvol -extpool p2 -cap 12 -name fb_0_#h -eam rotateexts 2200-2201
CMUC00025I mkfbvol: FB volume 2200 successfully created.
CMUC00025I mkfbvol: FB volume 2201 successfully created.
dsccli> chvolgrp -action add -volume 2200 v0
CMUC00031I chvolgrp: Volume group V0 successfully modified.
dsccli> rmfbvol -quiet -safe 2200-2201
CMUC00253E rmfbvol: Volume IBM.2107-75HAL91/2200 is assigned to a user-defined
volume group. No volumes were deleted.
dsccli> rmfbvol -quiet -safe 2201
CMUC00028I rmfbvol: FB volume 2201 successfully deleted.
dsccli> rmfbvol 2200
CMUC00027W rmfbvol: Are you sure you want to delete FB volume 2200? [Y/N]: y
CMUC00028I rmfbvol: FB volume 2200 successfully deleted.
```

Reinitializing space efficient FB volumes

Users can use the **initfbvol** command to reinitialize online space efficient FB volumes.

The command includes options to prevent the accidental reinitialization of volumes that are in use. An FB volume is considered to be in use if it is participating in a CS relationship or if the volume received any I/O operation in the previous 5 minutes. All data is lost when this command is used.

Volume reinitialization is controlled by the **-action releasespace** and **-force** parameters in the following manner (Example 10-43):

- ▶ The **releasespace** parameter is used to free all extents and tracks that are associated with the specified volume so they can then be reused by other space efficient volumes.
- ▶ The **-force** parameter is required in order for the **initfbvol** command to reinitialize the specified volume if it is in use.
- ▶ When both of these parameters are used, the user is prompted with a Y/N response in order for the command to proceed with the reinitialization process.

*Example 10-43 Using the **initfbvol** command to reinitialize a volume*

```

dscli> initfbvol -action releasespace -force 2005
CMUC00337W initfbvol: Are you sure that you want to free all extents and lose the
data associated with the FB volume 2005? [Y/N]:y
CMUC00340I initfbvol: 2005: The command releasespace has completed successfully.

```

10.3.6 Creating the volume groups

FB volumes are assigned to open system hosts by using volume groups. Do not confuse them with the *volume groups* term, which is used in AIX. An FB volume can be a member of multiple volume groups. Volumes can be added or removed from volume groups as required. Each volume group must be SCSI MAP256 or SCSI MASK, depending on the SCSI LUN address discovery method that is used by the OS to which the volume group is attached.

Note: There are two ways to create volume groups and map the volumes to the hosts. Volume groups can be created manually in single steps or automatically. The automatic method is done by using the **mkhost** and **chhost** commands, and it is the recommended method for mapping volumes to host systems.

Mapping of volumes to hosts or host groups is also called *storage partitioning*.

The following sections describe both ways, but understand that the manual way is only there for compatibility with earlier versions and must be applied with care to make sure that all the steps are done to fully reflect the results and views in the DS GUI.

Determining whether an open system host is SCSI MAP256 or SCSI MASK

Determine the type of SCSI host with which you are working. Then, run the **lshosttype** command with the **-type** parameter of **scsimask** and then **scsimap256**.

Example 10-44 shows the results of each command.

*Example 10-44 Listing the host types by running the **lshosttype** command*

```

dscli> lshosttype -type scsimask
HostType          Profile                                     AddrDiscovery  LBS
=====
Hp                HP - HP/UX                                 reportLUN      512
SVC               SAN Volume Controller                     reportLUN      512
SanFsAIX          IBM pSeries - AIX/SanFS                   reportLUN      512
pSeries           IBM pSeries - AIX                         reportLUN      512
pSeriesPowerswap IBM pSeries - AIX with Powerswap support  reportLUN      512
zLinux            IBM zSeries - zLinux                       reportLUN      512

```

```

dscli> lshosttype -type scsimap256
HostType  Profile                      AddrDiscovery LBS
=====
AppleOSX  Apple - OSX                    LUNPolling   512
Fujitsu   Fujitsu - Solaris              LUNPolling   512
HpTru64   HP - Tru64                     LUNPolling   512
HpVms     HP - Open VMS                  LUNPolling   512
Linux     Linux Server                   LUNPolling   512
Novell    Novell                          LUNPolling   512
SGI       SGI - IRIX                     LUNPolling   512
SanFsLinux - Linux/SanFS                 LUNPolling   512
Sun       SUN - Solaris                  LUNPolling   512
VMWare    VMWare                          LUNPolling   512
Windows   Windows Server                 LUNPolling   512
iLinux    IBM iSeries - iLinux           LUNPolling   512
nSeries   IBM N series Gateway           LUNPolling   512
pLinux    IBM pSeries - pLinux           LUNPolling   512

```

Creating a volume group

After you determine the host type, create a volume group. In Example 10-45, the example host type is AIX. In Example 10-44 on page 370, you can see that the address discovery method for AIX is scsimask.

Example 10-45 Creating a volume group by using mkvolgrp and displaying it by using lsvolgrp

```

dscli> mkvolgrp -type scsimask -volume 2000-2002,2100-2102 AIX_VG_01
CMUC00030I mkvolgrp: Volume group V1 successfully created.
dscli> lsvolgrp -l -type scsimask
Name          ID Type
=====
v0            V0 SCSI Mask
AIX_VG_01     V1 SCSI Mask
pE950_042     V2 SCSI Mask
pE950_048     V3 SCSI Mask
pseries_cluster V7 SCSI Mask
dscli> showvolgrp V1
Name AIX_VG_01
ID   V1
Type SCSI Mask
Vols 2000 2001 2002 2100 2101 2102

```

Adding or deleting volumes in a volume group

In this example, you add volumes 2000 - 2002 and 2100 - 2102 to the new volume group to evenly spread the workload across the two rank groups. You then list all available volume groups by using the **lsvolgrp** command. Finally, list the contents of volume group V1 because you created this volume group.

You might also want to add or remove volumes to this volume group later. To add or remove volumes, use the **chvolgrp** command with the **-action** parameter.

Example 10-46 shows adding volume 2003 to volume group V1, displaying the results, and then removing the volume.

Example 10-46 Changing a volume group by using chvolgrp

```
dsccli> chvolgrp -action add -volume 2003 V1
CMUC00031I chvolgrp: Volume group V1 successfully modified.
dsccli> showvolgrp V1
Name AIX_VG_01
ID V1
Type SCSI Mask
Vols 2000 2001 2002 2003 2100 2101 2102
dsccli> chvolgrp -action remove -volume 2003 V1
CMUC00031I chvolgrp: Volume group V1 successfully modified.
dsccli> showvolgrp V1
Name AIX_VG_01
ID V1
Type SCSI Mask
Vols 2000 2001 2002 2100 2101 2102
```

Important: Not all OSs can manage a volume removal. To determine the safest way to remove a volume from a host, see your OS documentation.

10.3.7 Creating host connections and clusters

The logical configuration process also includes creating host connections for your attached hosts. You must assign volume groups to those connections. Each host's host bus adapter (HBA) can be defined only once. Each host connection can have only one volume group that is assigned to it. A volume can be assigned to multiple volume groups.

You can use a set of cluster commands (**mkcluster**, **lscluster**, **showcluster**, and **rmcluster**) to create clusters to group hosts that have the same set of volume mappings, map or unmap volumes directly to these clusters, or both. These commands were added to provide consistency between the GUI and DS CLI. For more information, see "Creating open system clusters and hosts" on page 284.

Clusters are grouped hosts that must share volume access with each other. A cluster usually contains several hosts. Single hosts can exist without a cluster. Clusters are created by running the **mkcluster** command. This command does not need many parameters and is there only to organize hosts.

The **mkhost** command now has two generic host types that are available: Linux Server and Windows Server. These types were created to simplify and remove confusion when configuring these host types. You must define the host type first by running the **mkhost** command, as shown in Example 10-47.

Example 10-47 Creating generic host types Linux Server and Windows Server

```
dsccli> mkcluster cluster_1
CMUC00538I mkcluster: The cluster cluster_1 is successfully created.
Usage: mkhost [ { -help|-h|-? } ] [-v on|off] [-bnr on|off] [-dev storage_image_ID] -type
AIX|AIX with PowerSwap|HP OpenVMS|HP-UX|IBM i AS/400|iLinux|Linux RHEL|Linux SUSE|Linux
Server|N series Gateway|Novell|pLinux|SAN Volume Controller|Solaris|VMware|Windows
2003|Windows 2008|Windows 2012|Windows Server|zLinux [-hostport wwpn1[,wwpn2,...]]
[-cluster cluster_name] Host_Name | -
```

```

dscli> mkhost -type "Linux Server" -cluster cluster_1 Host_1
CMUC00530I mkhost: The host Host_1 is successfully created.
dscli> mkhost -type "Linux Server" -cluster cluster_1 Host_2
CMUC00530I mkhost: The host Host_2 is successfully created.
dscli> lshost
Name          Type          State  Cluster
=====
Host_1       Linux Server  Offline cluster_1
Host_2       Linux Server  Offline cluster_1
pE950_042    AIX           Online  -
pE950_048    AIX           Online  -
x3550_2_03   Windows Server Online  -
x3650_2_54   Linux Server  Online  -

```

More commands are also available: **chhost**, **lshost**, **showhost**, **chhost**, and **rmhost**. These commands were added to provide consistency between the DS GUI and DS CLI. Example 10-48 provides examples of the commands.

Example 10-48 Examples of the four extra CLI host commands

```

dscli> mkcluster cluster_2
CMUC00538I mkcluster: The cluster cluster_2 is successfully created.

dscli> chhost -cluster cluster_2 Host_1
CMUC00531I chhost: The host Host_1 is successfully modified.

dscli> lshost
Name          Type          State  Cluster
=====
Host_1       Linux Server  Offline cluster_2
Host_2       Linux Server  Offline cluster_1
pE950_042    AIX           Online  -
pE950_048    AIX           Offline -
x3550_2_03   Windows Server Online  -
x3650_2_54   Linux Server  Online  -

```

```

dscli> showhost Host_1
Name          Host_1
Type          Linux Server
State         Offline
AddrMode      scsimap256
Volumes       0
Host ports    0
I/O ports     16 (all)
AddrDiscovery lunpolling
LBS           512
Cluster       cluster_2

```

To link the logical hostname with a real physical connection, host ports must be assigned to the host by running the **mkhostport** command. This task also can be done by running the **mkhost** command with the **-hostport wwpn1[,wwpn2, ...]** option during host creation to save the additional configuration step.

To see a list of unassigned worldwide port names (WWPNs), which are already logged in to the storage system and represent the physical HBA ports of the hosts, run the **lshostport** command. Specifying **-unknown** shows all the free ports and **-login** shows all the logged-in ports. It takes a while until the storage system shows the new logged-in ports. It is also possible to add them manually, and they do not need to be logged in to create the host configuration and volume mappings.

Example 10-49 shows the host port connection.

Example 10-49 Host port connection

```

dsccli> lshostport -login
WWNN                WWPNN                Host                I/O ports ESSI0port
=====
20000000C9F607E8  10000000C9F607E8  -                    -            I0032
20000000C9F607E9  10000000C9F607E9  -                    -            I0103
2000001B329CBB21  2100001B329CBB21  x3650_2_54          16 (all)    I0032
20000024FF2D0F86  21000024FF2D0F86  x3550_2_03          16 (all)    I0103
20000024FF2D0F87  21000024FF2D0F87  x3550_2_03          16 (all)    I0032
20000024FF41D1CE  21000024FF41D1CE  x3550_2_03          16 (all)    I0103
20000024FF41D1CF  21000024FF41D1CF  x3550_2_03          16 (all)    I0032
20000120FA13B914  10000090FA13B914  -                    -            I0032
20000120FA13B915  10000090FA13B915  -                    -            I0103
2001001B32BCBB21  2101001B32BCBB21  x3650_2_54          16 (all)    I0103
C05076068D1B004A  C05076068D1B004A  pE950_048           16 (all)    I0032
C05076068D1B004C  C05076068D1B004C  pE950_048           16 (all)    I0103
dsccli> lshostport -unknown
Date/Time: June 27, 2022 5:50:08 PM CEST IBM DSCLI Version: 7.9.30.154 DS: IBM.2107-75HAL91
WWNN                WWPNN                ESSI0port
=====
20000000C9F607E8  10000000C9F607E8  I0032
20000000C9F607E9  10000000C9F607E9  I0103
20000120FA13B914  10000090FA13B914  I0032
20000120FA13B915  10000090FA13B915  I0103
dsccli> mkhostport -host Host_1 "10000000C9F607E8"
CMUC00542I mkhostport: The host port 10000000C9F607E8 is successfully created.
dsccli> mkhostport -host Host_1 "10000000C9F607E9"
CMUC00542I mkhostport: The host port 10000000C9F607E9 is successfully created.
dsccli> lshostport
WWPN                Name                Type                State
=====
10000000C9F607E8  Host_1              Linux Server        logged_in
10000000C9F607E9  Host_1              Linux Server        logged_in
2100001B329CBB21  x3650_2_54          Linux Server        logged_in
21000024FF2D0F86  x3550_2_03          Windows Server      logged_in
21000024FF2D0F87  x3550_2_03          Windows Server      logged_in
21000024FF41D1CE  x3550_2_03          Windows Server      logged_in
21000024FF41D1CF  x3550_2_03          Windows Server      logged_in
2101001B32BCBB21  x3650_2_54          Linux Server        logged_in
C05076068D1B004A  pE950_048           AIX                  logged_in
C05076068D1B004C  pE950_048           AIX                  logged_in
dsccli> lshost -l
Name                Type                State  Cluster  AddrMode  Volumes  Host ports  I/O ports
=====
Host_1              Linux Server        Online cluster_2  scsimap256  0          2 16 (all)
Host_2              Linux Server        Offline cluster_1  scsimap256  0          0 16 (all)
pE950_042           AIX                  Offline -          scsimask     0          0 16 (all)
pE950_048           AIX                  Online -          scsimask     0          2 16 (all)
x3550_2_03          Windows Server      Online -          scsimap256  6          4 16 (all)
x3650_2_54          Linux Server        Online -          scsimap256  8          2 16 (all)

```

The last step is to assign volumes to the host or cluster to allow host access to the volumes. The volumes that are assigned to the host are seen only by the host, and the volumes that are assigned to the cluster can be seen by all hosts inside the cluster. The volume groups for the cluster and the host are generated automatically by running the **chhost/chcluster -action map** command. The automatically created volume groups have the same name as the host or cluster itself but are different objects.

Example 10-50 shows the steps for a cluster and a host.

Example 10-50 Volume to host mapping

```

dscli> lsvolgrp -type scsimap256
Name          ID Type
=====
x3650_2_54 V4 SCSI Map 256
x3550_2_03 V8 SCSI Map 256
dscli> chhost -action map -volume 2000-2003 Host_1
CMUC00531I chhost: The host Host_1 is successfully modified.
dscli> lsvolgrp -type scsimap256
Name          ID Type
=====
x3650_2_54 V4 SCSI Map 256
Host_1      V5 SCSI Map 256
x3550_2_03 V8 SCSI Map 256
dscli> chcluster -action map -volume 2100-2103 cluster_2
CMUC00541I chcluster: The cluster cluster_2 is successfully modified.
dscli> lsvolgrp -type scsimap256
Name          ID Type
=====
x3650_2_54 V4 SCSI Map 256
Host_1      V5 SCSI Map 256
cluster_2  V6 SCSI Map 256
x3550_2_03 V8 SCSI Map 256
dscli> lshost -l
Name          Type          State  Cluster  AddrMode  Volumes  Host ports  I/O ports
=====
Host_1      Linux Server  Online cluster_2 scsimap256  8         2 16 (all)
Host_2       Linux Server  Offline cluster_1 scsimap256  0         0 16 (all)
pE950_042    AIX           Offline -          scsimask   0         0 16 (all)
pE950_048    AIX           Online  -          scsimask   0         2 16 (all)
x3550_2_03   Windows Server Online -          scsimap256 6         4 16 (all)
x3650_2_54   Linux Server  Online  -          scsimap256 8         2 16 (all)

```

Example 10-51 shows some change and removal commands. Unassigning the host from the cluster means that the host also keeps the cluster volume mappings and the cluster also keeps them.

The automatically created volumes groups are removed only if the host is removed or you run the **rmvolgrp** command.

Example 10-51 Host change and removal

```

dscli> showhost Host_1
Name          Host_1
Type          Linux Server
State         Online
AddrMode      scsimap256
Volumes       8
Host ports    2
I/O ports     16 (all)

```

```

AddrDiscovery lunpolling
LBS          512
Cluster      cluster_2
dscli> chhost -cluster cluster_2 Host_2
CMUC00531I chhost: The host Host_2 is successfully modified.
dscli> chhost -unassign Host_1
CMUC00531I chhost: The host Host_1 is successfully modified.

dscli> showhost Host_1
Name          Host_1
Type          Linux Server
State         Online
AddrMode      scsimap256
Volumes      8
Host ports    2
I/O ports     16 (all)
AddrDiscovery lunpolling
LBS           512
Cluster      -

dscli> rmhost -quiet Host_1
CMUC00533E rmhost: The host IBM.2107-75HAL91/Host_1 cannot be removed because it contains
volumes.
dscli> lsvolgrp -type scsimap256
Name          ID Type
=====
x3650_2_54 V4 SCSI Map 256
Host_1      V5 SCSI Map 256
cluster_2    V6 SCSI Map 256
x3550_2_03 V8 SCSI Map 256
Host_2       V9 SCSI Map 256
dscli> showvolgrp v5
Name Host_1
ID    V5
Type SCSI Map 256
Vols 2000 2001 2002 2003 2101 2100 2102 2103
dscli> chvolgrp -action remove -volume 2000-2003,2100-2103 v5
CMUC00031I chvolgrp: Volume group V9 successfully modified.
dscli> rmhost Host_1
CMUC00529W rmhost: Are you sure that you want to delete host Host_1? [Y/N]: y
CMUC00528I rmhost: The host Host_1 is successfully deleted.
dscli> lsvolgrp -type scsimap256
Name          ID Type
=====
x3650_2_54 V4 SCSI Map 256
cluster_2    V6 SCSI Map 256
x3550_2_03 V8 SCSI Map 256
Host_2       V9 SCSI Map 256

```

Connecting a host to an existing volume group

Example 10-45 on page 371 shows the already created volume group for an example AIX host connection. The volumes are already assigned to the volume group.

Example 10-52 on page 377 shows the creation of a host connection that represents two HBA ports in this AIX host. Use the **-hosttype** parameter to include the host type that you used in Example 10-44 on page 370. Allocate it to volume group V1. If the storage area network (SAN) zoning is correct, the host can see the LUNs in volume group V1.

Example 10-52 Creating host connections by using mkhostconnect and lshostconnect

```
dsccli> lshostport -unknown
WWNN                WWPN                ESSIOport
=====
20000120FA13B914  10000090FA13B914  I0032
20000120FA13B915  10000090FA13B915  I0103
dsccli> lsvolgrp -type scsimask
Name                ID Type
=====
v0                  V0 SCSI Mask
AIX_VG_01         V1 SCSI Mask
pE950_042          V2 SCSI Mask
pE950_048          V3 SCSI Mask
pseries_cluster    V7 SCSI Mask
dsccli> showvolgrp V1
Name AIX_VG_01
ID   V1
Type SCSI Mask
Vols 2000 2001 2002 2100 2101 2102
dsccli> mkhostconnect -wwname 10000090FA13B914 -hosttype pSeries -volgrp V1 AIX_Server_01
CMUC00012I mkhostconnect: Host connection 000A successfully created.
dsccli> mkhostconnect -wwname 10000090FA13B915 -hosttype pSeries -volgrp V1 AIX_Server_01
CMUC00012I mkhostconnect: Host connection 000B successfully created.
dsccli> lshostconnect
Name                ID WWPN                HostType Profile                portgrp volgrpID ESSIOport
=====
AIX_Server_01 000A 10000090FA13B914 pSeries  IBM pSeries - AIX                0 V1      all
AIX_Server_01 000B 10000090FA13B915 pSeries  IBM pSeries - AIX                0 V1      all
```

You can also use **-profile** instead of **-hosttype**. However, this method is not a best practice. Using the **-hosttype** parameter reflects both parameters (**-profile** and **-hosttype**). In contrast, using **-profile** leaves the **-hosttype** column unpopulated.

The option in the **mkhostconnect** command to restrict access only to certain I/O ports is also available by using the **-ioport** parameter. Restricting access in this way is unnecessary. If you want to restrict access for certain hosts to certain I/O ports on the DS8900F, perform zoning on your SAN switch.

The **mkhostconnect** command normally is sufficient to allow the volumes to access the specified host ports. The command works, but it is not reflected in the modernized GUI interface. The modernized GUI interface introduced host and cluster grouping for easier management of groups of hosts with many host ports. If no host or cluster is assigned to the created connection, the GUI still shows the ports as unassigned host ports with mapped volumes.

Figure 10-8 shows the results in the DS GUI.

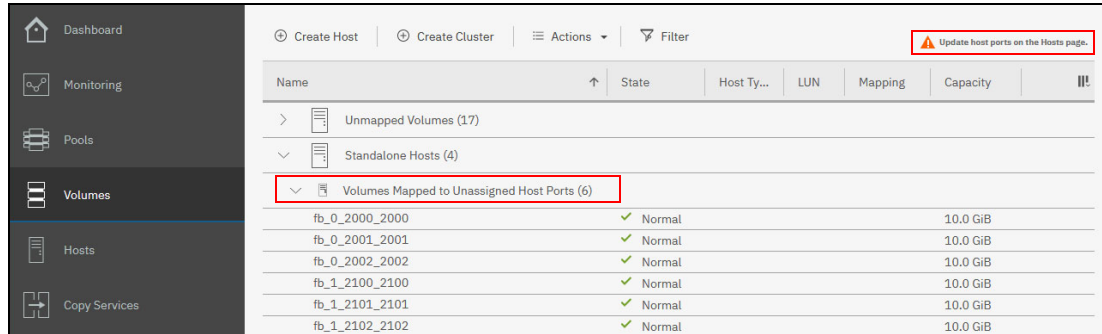


Figure 10-8 Volumes that are mapped to a host port without a host

The `lshostconnect -l` command in Example 10-53 shows that the relationship between the volume group and a host connection was not built up. The assigned host is missing in the last column and portgrp 0 is used, which is not recommended because it is the default port group for new host ports. There is no host that is created yet for the AIX connection in our example.

The first column does not show the hostname: It is a symbolic name for the connection for better recognition. The ID field makes the connection unique.

Example 10-53 The lshostconnect -l command for full view

```
dscli> lshostconnect -l
```

Name	ID	WWPN	HostType	LBS	addrDiscovery	Profile	portgrp	volgrpID	ESSIOport	host
x3650_2_54	0001	2100001B329CBB21	Linux	512	LUNPolling	Linux Server	1 V4		all	x3650_2_54
x3650_2_54	0002	2101001B32BCBB21	Linux	512	LUNPolling	Linux Server	1 V4		all	x3650_2_54
x3550_2_03	0004	21000024FF2D0F86	Windows	512	LUNPolling	Windows Server	2 V8		all	x3550_2_03
x3550_2_03	0005	21000024FF2D0F87	Windows	512	LUNPolling	Windows Server	2 V8		all	x3550_2_03
x3550_2_03	0006	21000024FF41D1CE	Windows	512	LUNPolling	Windows Server	2 V8		all	x3550_2_03
x3550_2_03	0007	21000024FF41D1CF	Windows	512	LUNPolling	Windows Server	2 V8		all	x3550_2_03
AIX_Server_01	000A	10000090FA13B914	pSeries	512	reportLUN	IBM pSeries - AIX	0 V1		all	-
AIX_Server_01	000B	10000090FA13B915	pSeries	512	reportLUN	IBM pSeries - AIX	0 V1		all	-
pE950_048	000F	C05076068D1B004A	pSeries	512	reportLUN	IBM pSeries - AIX	4 V3		all	pE950_048
pE950_048	0010	C05076068D1B004C	pSeries	512	reportLUN	IBM pSeries - AIX	4 V3		all	pE950_048

Note: As a best practice, it is not advisable to use the host port group ID of 0. This ID ties together a group of SCSI host port objects that are accessing a common volume group. If the port group value is set to zero, the host port is not associated with any port group. It is used by default for ports that are not grouped yet.

When hosts are created, you can specify the `-portgrp` parameter. By using a unique port group number for each attached server, you can detect servers with multiple HBAs.

If you want to use a single command to change the assigned volume group of several host connections at the same time, assign the host connections to a unique port group. Then, run the `managehostconnect` command. This command changes the assigned volume group for all host connections that are assigned to a particular port group.

Changing host connections

If you want to change a host connection, run the **chhostconnect** command. This command can be used to change nearly all parameters of the host connection, except for the WWPN.

Example 10-54 shows the steps to finish the configuration of the mapping. A host must be created, and then the new host must be assigned to the existing connections A and B and volume group V1 relationship.

Example 10-54 Host and volume group relationship

```
dscli> mkhost -type AIX Host_AIX
CMUC00530I mkhost: The host Host_AIX is successfully created.
dscli> lshost -l
Name          Type          State  Cluster  AddrMode  Volumes  Host ports  I/O ports
=====
Host_AIX     AIX              Offline -          scsimask    0          0 16 (all)
...
dscli> lshostport -l
WWPN          Name          Type          State  Volumes  AddrMode  I/O ports  Host
=====
10000090FA13B914 AIX_Server_01 AIX              logged_in    6    scsimask 16 (all) -
10000090FA13B915 AIX_Server_01 AIX              logged_in    6    scsimask 16 (all) -
...
dscli> lshostconnect
Name          ID  WWPN          HostType Profile          portgrp volgrpID ESSIOport
=====
...
AIX_Server_01 000A 10000090FA13B914 pSeries  IBM pSeries - AIX    0 V1    all
AIX_Server_01 000B 10000090FA13B915 pSeries  IBM pSeries - AIX    0 V1    all
...
dscli> chhostconnect -host Host_AIX A
CMUC00527W chhostconnect: The volumes will be merged into the host. Are you sure that you
want to modify host connection A? [Y/N]: y
CMUC00013I chhostconnect: Host connection 000A successfully modified.
dscli> chhostconnect -quiet -host Host_AIX B
CMUC00013I chhostconnect: Host connection 000B successfully modified.
dscli> lshost -l
Name          Type          State  Cluster  AddrMode  Volumes  Host ports  I/O ports
=====
Host_AIX     AIX              Online -          scsimask    6          2 16 (all)
...
dscli> lshostport -l
WWPN          Name          Type          State  Volumes  AddrMode  I/O ports  Host
=====
10000090FA13B914 AIX_Server_01 AIX              logged_in    6    scsimask 16 (all) Host_AIX
10000090FA13B915 AIX_Server_01 AIX              logged_in    6    scsimask 16 (all) Host_AIX
...

```

After completing the steps, the configuration is done.

Example 10-55 with the command `lshostconnect -l` shows that the relationship between host, connection, and volume group is made. The `chhostconnect` command automatically changed the host port group to a value different from zero.

Example 10-55 The lshostconnect -l command

```

dscli> lshostconnect -l
Name          ID    WWPN          HostType LBS addrDiscovery Profile          portgrp volgrpID ESSIOport host
=====
...
AIX_Server_01 000A 10000090FA13B914 pSeries 512 reportLUN    IBM pSeries - AIX      5 V1    all    Host_AIX
AIX_Server_01 000B 10000090FA13B915 pSeries 512 reportLUN    IBM pSeries - AIX      5 V1    all    Host_AIX
...

```

Note: Using the `mkvolgrp` and `mkhostconnect` commands for storage partitioning to map volumes to hosts is not the preferred method. It is available for compatibility for earlier and existing volume groups. It is better to use the `mkhost` command from the beginning to assign hosts, host ports, volume groups, and volume mappings together. It combines the needed functions in one command and makes sure that no step is forgotten. It also reduces the number of steps that is needed.

10.3.8 Mapping open system host disks to storage unit volumes

When you assign volumes to an open system host and install the DS CLI on a host, you can run the `lshostvol` DS CLI command on that host. This command maps assigned LUNs to open system host volume names.

You log on to this host and start DS CLI. It does not matter which HMC that you connect to when you use the DS CLI. Then, run the `lshostvol` command.

Important: The `lshostvol` command communicates only with the OS of the host on which the DS CLI is installed. You cannot run this command on one host to see the attached disks of another host.

Note: The Subsystem Device Driver Path Control Module (SDDPCM) (a multipath solution on AIX) and Subsystem Device Driver Device Specific Module (SDDDSM) (a multipath solution on Windows) are no longer developed for DS8000. Instead, use an OS-native solution such as AIX Multipath I/O Path Control Module (AIXPCM) (the AIX default multipath solution) or Microsoft Device Specific Module (MSDSM) (the Windows default multipath solution). These solutions are fully supported on open systems.

10.4 DS8900F storage configuration for the CKD volumes

This list contains the steps to configure CKD storage in the DS8880:

1. Create the arrays.
2. Create the CKD ranks.
3. Create the CKD extent pools.
4. Create the logical control units (LCUs).
5. Create the CKD volumes.

You do not need to create volume groups or host connects for CKD volumes. If I/O ports in FICON mode exist, access to CKD volumes by FICON hosts is granted automatically, and follows the specifications in the input/output definition file (IODF).

10.4.1 Creating the arrays

Array creation for CKD volumes is the same as for FB volumes. For more information, see 10.3.2, “Creating the arrays” on page 358.

10.4.2 Creating the ranks and extent pools

When ranks and extent pools are created, you must specify **-stgtype ckd**. Then, you can create the extent pool, as shown in Example 10-56.

Example 10-56 Rank and extent pool creation for CKD volumes

```
dscli> mkrank -array A2 -stgtype ckd
CMUC00007I mkrank: Rank R2 successfully created.
dscli> lsrank
ID Group State      datastate Array RAIDtype extpoolID stgtype
=====
R0    0 Normal    Normal    A0          6 P0       fb
R1    1 Normal    Normal    A1          6 P1       fb
R2    - Unassigned Normal    A2          6 -        ckd
dscli> mkextpool -rankgrp 0 -stgtype ckd CKD_HPF_0
CMUC00000I mkextpool: Extent Pool P2 successfully created.
dscli> chrnk -extpool P2 R2
CMUC00008I chrnk: Rank R2 successfully modified.
dscli> lsextpool
Name      ID stgtype rankgrp status availstor (2^30B) %allocated available reserved numvols
=====
ITSO_FB  P0 fb          0 below      15232         12  15232         1    17
ITSO_FB  P1 fb          1 below      15833         8   15833         1    7
CKD_HPF_0 P2 ckd         0 below      7135          0   8098         1    0
```

When defining a rank, you can also specify the extent size. You can have ranks and extent pools with large 1113 cylinder CKD extents, or small 21 cylinder CKD extents. The extent unit is specified with the **-extsize** parameter of the **mkrank** command. The first rank that is added to an extent pool determines the extent size of the extent pool. Example 10-57 shows ranks with different extent sizes.

Example 10-57 CKD ranks with different extent sizes

```
dscli> lsrank -l
Date/Time: June 27, 2022 5:55:29 PM CEST IBM DSCLI Version: 7.9.30.154 DS: IBM.2107-75HAL91
ID Group State datastate Array RAIDtype extpoolID extpoolnam stgtype exts usedexts keygrp marray extsize (cap)
=====
R0    0 Normal Normal    A0          6 P0      ITSO_FB  fb    17338    2105    1 MA4    1GiB
R1    1 Normal Normal    A1          6 P1      ITSO_FB  fb    17338    1504    1 MA3    1GiB
R2    0 Normal Normal    A2          6 P2      CKD_HPF_0 ckd    8099     0       1 MA1    1113cyl
R3    1 Normal Normal    A3          6 P3      CKD_HPF_1 ckd   429258  0       1 MA2    21cyl
```

10.4.3 Logical control unit creation

In a CKD environment, you must create LCUs before the volumes are created. In Example 10-58, you can see what happens if you try to create a CKD volume without creating an LCU first.

Example 10-58 Trying to create CKD volumes without creating an LCU first

```
dsccli> mkckdvol -extpool p2 -cap 262668 -name CKD_EAV1_#h C200
CMUN02282E mkckdvol: C200: Unable to create CKD logical volume: CKD volumes require a CKD
logical subsystem.
```

To create the LCUs, run the **mk1cu** command. The command uses the following format:

```
mk1cu -qty XX -id XX -ss XXXX
```

Note: For the z/OS hardware definition, the subsystem identifier (SSID) (-ss -id) must be unique for all connected storage systems. The z/OS hardware admin usually provides the SSID to use.

To display the LCUs that you created, run the **ls1cu** command.

Example 10-59 shows the creation of two LCUs by running the **mk1cu** command and then listing the created LCUs by running the **ls1cu** command. By default, the LCUs that were created are the 3990-6 type.

Example 10-59 Creating a logical control unit by running mk1cu

```
dsccli> mk1cu -qty 2 -id BC -ss 91BC
CMUC00017I mk1cu: LCU BC successfully created.
CMUC00017I mk1cu: LCU BD successfully created.
dsccli> ls1cu
ID Group addrgrp confgvols subsys conbasetype
=====
BC      0 B           0 0x91BC 3990-6
BD      1 B           0 0x91BD 3990-6
```

Because two LCUs were created by using the parameter **-qty 2**, the first LCU, which is ID BC (an even number), is in address group 0, which equates to rank group 0. The second LCU, which is ID BD (an odd number), is in address group 1, which equates to rank group 1. By placing the LCUs into both address groups, performance is maximized by spreading the workload across both servers in the DS8900F.

Important: For the DS8900F, the CKD LCUs can be ID 00 - FE. The LCUs fit into one of 16 address groups. Address group 0 is LCUs 00 - 0F, address group 1 is LCUs 10 - 1F, and so on, except group F is F0 - FE. If you create a CKD LCU in an address group, that address group cannot be used for FB volumes. Likewise, if, for example, FB volumes were in LSS 40 - 4F (address group 4), that address group cannot be used for CKD. Be aware of this limitation when you plan the volume layout in a mixed FB and CKD DS8900F. Each LCU can manage a maximum of 256 volumes, including alias volumes for the parallel access volume (PAV) feature.

10.4.4 Creating the CKD volumes

Now that an LCU is created, the CKD volumes can be created by using the `mkckdvol` command. The `mkckdvol` command uses the following format:

```
mkckdvol -extpool P2 -cap 262668 -datatype 3390-A -eam rotatevols -name
CKD_EAV1_#h BC06
```

The greatest difference with CKD volumes is that the capacity is expressed in cylinders or as mod1 (Model 1) extents (1113 cylinders). To not waste space, use volume capacities that are a multiple of 1113 cylinders.

The support for extended address volumes (EAVs) was enhanced. The DS8900F supports EAV volumes up to 1,182,006 cylinders. The EAV device type is called *3390 Model A*.

Important: For 3390-A volumes, the size can be specified as 1 - 65,520 in increments of 1, and from 65,667, which is the next multiple of 1113, to 1,182,006 in increments of 1113.

The last parameter in the command is the `volume_ID`. This value determines the LCU that the volume belongs to and the unit address (UA) for the volume. Both of these values must be matched to a control unit and device definition in the input/output configuration data set (IOCDs) that an IBM Z system server uses to access the volume.

The `volume_ID` has a format of *LLVV*. *LL* (00 - FE) equals the LCU to which the volume belongs, and *VV* (00 - FF) equals the offset for the volume. Only one volume of an LCU can use a unique *VV* of 00 - FF.

Example 10-60 shows the creation of 3390-A volumes with a capacity of 262,668 cylinders that are assigned to LCU BC with an offset of 00 - 05.

Example 10-60 Creating CKD volumes by using `mkckdvol`

```
dsccli> mkckdvol -extpool P2 -cap 262668 -datatype 3390-A -eam rotatevols -name CKD_EAV1_#h BC00-BC05
CMUC00021I mkckdvol: CKD Volume BC00 successfully created.
CMUC00021I mkckdvol: CKD Volume BC01 successfully created.
CMUC00021I mkckdvol: CKD Volume BC02 successfully created.
CMUC00021I mkckdvol: CKD Volume BC03 successfully created.
CMUC00021I mkckdvol: CKD Volume BC04 successfully created.
CMUC00021I mkckdvol: CKD Volume BC05 successfully created.
dsccli> lsckdvol
Name          ID  accstate  datastate  configstate  deviceMTM  voltype  orgbvols  extpool  cap (cyl)
=====
CKD_EAV1_BC00 BC00 Online    Normal    Normal     3390-A     CKD Base -    P2      262668
CKD_EAV1_BC01 BC01 Online    Normal    Normal     3390-A     CKD Base -    P2      262668
CKD_EAV1_BC02 BC02 Online    Normal    Normal     3390-A     CKD Base -    P2      262668
CKD_EAV1_BC03 BC03 Online    Normal    Normal     3390-A     CKD Base -    P2      262668
CKD_EAV1_BC04 BC04 Online    Normal    Normal     3390-A     CKD Base -    P2      262668
CKD_EAV1_BC05 BC05 Online    Normal    Normal     3390-A     CKD Base -    P2      262668
```

You can create only CKD volumes in LCUs that you already created. Volumes in even-numbered LCUs must be created from an extent pool that belongs to rank group 0. Volumes in odd-numbered LCUs must be created from an extent pool in rank group 1. With one `mkckdvol` command, volumes for one LCU can be defined.

Important: You can configure a volume to belong to a certain resource group by using the `-resgrp <RG_ID>` flag in the `mkckdvol` command. For more information, see *IBM System Storage DS8000 Copy Services Scope Management and Resource Groups*, REDP-4758.

Defining thin-provisioned extent space efficient CKD volumes

The DS8900F supports thin-provisioned volumes. A thin-provisioned volume is created by the DS CLI by specifying the SAM of ESE for a volume. This process is done with the **-sam ese** option of the **mkckdvol** command, as shown in Example 10-61.

Example 10-61 Creating a thin-provisioned CKD volume

```
dsccli> mkckdvol -extpool p2 -cap 1000 -sam ese BC06
CMUC00021I mkckdvol: CKD Volume BC06 successfully created.
```

Space release of thin-provisioned ESE CKD volumes

It is possible to release space at an extent level for CKD volumes. This feature is enabled by the z/OS utility **DFSMSdss**, which performs the release operation. **DFSMSdss** includes a parameter, **SPACERe1**, with its options. A corresponding IBM RACF FACILITY class profile provides protection as well. It is available on z/OS V2.1 and z/OS V2.2 with Program Temporary Fix (PTF) OA50675.

More DS CLI commands are available to control and protect the space in an extent pool for thin-provisioned volumes. One of these commands, the **mksestg** command, reserves space for thin-provisioned volumes. For more information about thin-provisioning, see *IBM DS8880 Thin Provisioning (Updated for Release 8.5)*, REDP-5343.

Storage pool striping

When a volume is created, you can choose how the volume is allocated in an extent pool with several ranks. The extents of a volume can be kept together in one rank (if enough free space exists on that rank). The next rank is used when the next volume is created. This allocation method is called *rotate volumes*.

You can also specify that you want the extents of the volume to be evenly distributed across all ranks within the extent pool. This allocation method is called *rotate extents*.

Rotate extents: For the DS8900F, the default allocation policy is rotate extents.

The EAM is specified with the **-eam rotateexts** or the **-eam rotatevols** option of the **mkckdvol** command (Example 10-62).

Example 10-62 Creating a CKD volume with extent pool striping

```
dsccli> mkckdvol -extpool p3 -cap 10017 -name CKD_EAV1_#h -eam rotateexts BD00-BD03
CMUC00021I mkckdvol: CKD Volume BD00 successfully created.
CMUC00021I mkckdvol: CKD Volume BD01 successfully created.
CMUC00021I mkckdvol: CKD Volume BD02 successfully created.
CMUC00021I mkckdvol: CKD Volume BD03 successfully created.
```

The **showckdvol** command with the **-rank** option (Example 10-63) shows that the volume that was created is distributed across two ranks. It also displays how many extents on each rank were allocated for this volume. In this example, the pool P3 uses small 21-cylinder CKD extents.

Example 10-63 Obtaining information about a striped CKD volume

```
dsccli> showckdvol -rank BD00
Name                CKD_EAV1_BD00
ID                  BD00
accstate            Online
```

```

datastate          Normal
configstate       Normal
deviceMTM         3390-9
volser            -
datatype          3390
voltype           CKD Base
orgbvols          -
addrgrp          B
extpool           P3
exts              477
cap (cyl)         10017
cap (10^9B)      8.5
cap (2^30B)      7.9
ranks             2
sam               Standard
repcapalloc       -
eam               rotateexts
reqcap (cyl)     10017
cap (Mod1)       9.0
realextents      477
virtualextents  2
realcap (cyl)   10017
migrating         0
migratingcap (cyl) 0
perfgrp          PG0
migratingfrom    -
resgrp           RG0
tierassignstatus Unknown
tierassignerror  -
tierassignorder  Unknown
tierassigntarget Unknown
%tierassigned    0
etmonpauseremain -
etmonitorreset   unknown
safeguardedcap (cyl) -
safeguardedloc  -
usedsafeguardedcap (cyl) -
safeguarded      no
SGC Recovered    no
=====Rank extents=====
rank extents capacity (MiB/cyl) metadata
=====
R2      239 5019          no
R3      238 4998          yes

```

Dynamic Volume Expansion

A volume can be expanded without removing the data within the volume. You can specify a new capacity by using the **chckdvol** command, as shown in Example 10-64. The new capacity must be larger than the previous capacity. *You cannot shrink the volume.*

Example 10-64 Expanding a striped CKD volume

```

dscli> chckdvol -cap 30051 BD00
CMUC00332W chckdvol: Some host operating systems do not support changing the volume size.
Data can be at risk if the host does not support this action. Are you sure that you want to
resize the volume? [Y/N]: y
CMUC00022I chckdvol: CKD Volume BD00 successfully modified.

```

Because the original volume used the `rotateexts` attribute, the additional extents are also striped, as shown in Example 10-65. In this example, the pool P3 is using small 21-cylinder CKD extents.

Example 10-65 Checking the status of an expanded CKD volume

```

dscli> showckdvol -rank BD00
Name                CKD_EAV1_BD00
ID                  BD00
accstate            Online
datastate           Normal
configstate         Normal
deviceMTM           3390-9
volser              -
datatype            3390
voltype             CKD Base
orgbvols            -
addrgrp             B
extpool             P3
exts                1431
cap (cyl)           30051
cap (10^9B)         25.5
cap (2^30B)         23.8
ranks               2
sam                 Standard
repcapalloc         -
eam                 rotateexts
reqcap (cyl)        30051
cap (Mod1)          27.0
realextents         1431
virtualextents     4
realcap (cyl)       30051
migrating           0
migratingcap (cyl) 0
perfgrp            PGO
migratingfrom       -
resgrp             RGO
tierassignstatus    Unknown
tierassignerror     -
tierassignorder     Unknown
tierassigntarget    Unknown
%tierassigned       0
etmonpauseremain    -
etmonitorreset      unknown
safeguardedcap (cyl) -
safeguardedloc      -
usedsafeguardedcap (cyl) -
safeguarded         no
SGC Recovered       no
=====Rank extents=====
rank extents capacity (MiB/cyl) metadata
=====
R2      716 15036          yes
R3      715 15015          yes

```

Important: Before you can expand a volume, you first must delete all CS relationships for that volume. Also, you cannot specify both `-cap` and `-datatype` in the same `chckdvol` command.

It is possible to expand a 3390 Model 9 volume to a 3390 Model A. Expand the volume by specifying new capacity for an existing Model 9 volume. When you increase the size of a 3390-9 volume beyond 65,520 cylinders, its device type automatically changes to 3390-A.

Important: A 3390 Model A can be used only in z/OS V1.10 (depending on the size of the volume) and later, as shown in Example 10-66.

Example 10-66 Expanding a 3390 to a 3390-A

*** Command to show CKD volume definition before expansion:

```

dscli> showckdvol -rank BD01
Name                CKD_EAV1_BD01
ID                  BD01
accstate            Online
datastate           Normal
configstate         Normal
deviceMTM           3390-9
volser              -
datatype            3390
voltype             CKD Base
orgbvols            -
addrgrp             B
extpool             P3
exts                477
cap (cyl)           10017
cap (10^9B)         8.5
cap (2^30B)         7.9
ranks               2
sam                 Standard
repcapalloc         -
eam                 rotateexts
reqcap (cyl)        10017
cap (Mod1)          9.0
realextensts        477
virtualextents     2
realcap (cyl)       10017
migrating           0
migratingcap (cyl) 0
perfgrp            PGO
migratingfrom       -
resgrp              RGO
tierassignstatus    Unknown
tierassignerror     -
tierassignorder     Unknown
tierassigntarget    Unknown
%tierassigned       0
etmonpauseremain   -
etmonitorreset      unknown
safeguardedcap (cyl) -
safeguardedloc      -
usedsafeguardedcap (cyl) -
safeguarded         no
SGC Recovered       no
=====Rank extents=====
rank extents capacity (MiB/cyl) metadata
=====
R2      239 5019          no
R3      238 4998          yes

```

*** Command to expand CKD volume from 3390-9 to 3390-A:

```
dscli> chckdvol -cap 262668 BD01
```

```
CMUC00332W chckdvol: Some host operating systems do not support changing the volume size.
Data can be at risk if the host does not support this action. Are you sure that you want to
resize the volume? [Y/N]: y
CMUC00022I chckdvol: CKD Volume BD01 successfully modified.
```

*** Command to show CKD volume definition after expansion:

```
dscli> showckdvol -rank BD01
```

```
Name          CKD_EAV1_BD01
ID            BD01
accstate      Online
datastate     Normal
configstate   Normal
deviceMTM     3390-A
volser        -
datatype      3390-A
voltype       CKD Base
orgbvols      -
addrgrp       B
extpool       P3
exts          12508
cap (cyl)     262668
cap (10^9B)   223.3
cap (2^30B)   207.9
ranks         2
sam           Standard
repcapalloc   -
eam           rotateexts
reqcap (cyl)  262668
cap (Mod1)    236.0
realextents   12508
virtualextents 25
realcap (cyl) 262668
migrating     0
migratingcap (cyl) 0
perfgrp       PG0
migratingfrom -
resgrp        RG0
tierassignstatus Unknown
tierassignerror -
tierassignorder Unknown
tierassigntarget Unknown
%tierassigned 0
etmonpauseremain -
etmonitorreset unknown
safeguardedcap (cyl) -
safeguardedloc -
usedsafeguardedcap (cyl) -
safeguarded   no
SGC Recovered no
=====Rank extents=====
rank extents capacity (MiB/cyl) metadata
=====
R2      6254 131334          yes
R3      6254 131334          yes
```

You *cannot* reduce the size of a volume. If you try to reduce the size, an error message is displayed.

CKD volumes can be deleted by using the **rmckdvol** command. FB volumes can be deleted by using the **rmfbvol** command.

The command includes a capability to prevent the accidental deletion of volumes that are in use. A CKD volume is considered *in use* if it participates in a CS relationship, or if the IBM Z path mask indicates that the volume is in a grouped state or online to any host system.

If the **-force** parameter is not specified with the command, volumes that are in use are not deleted. If multiple volumes are specified and several volumes are in use and several volumes are not, the volumes that are not in use are deleted.

If the **-force** parameter is specified on the command, the volumes are deleted without checking to see whether they are in use.

Example 10-67 shows an attempt to delete two volumes, BD02 and BD03. Volume BD02 is online on a host. Volume BD03 is not online on any host and not in a CS relationship. The **rmckdvol BD02-BD03** command deletes only volume BD03, which is offline. To delete volume BD02, use the **-force** parameter.

Example 10-67 Deleting CKD volumes

```
dsccli> lsckdvol BD02-BD03
Name          ID  accstate  datastate  configstate  deviceMTM  voltype  orgbvols  extpool  cap (cyl)
=====
CKD_EAV1_BD02 BD02 Online    Normal    Normal     3390-9     CKD Base -      P3      10017
CKD_EAV1_BD03 BD03 Online    Normal    Normal     3390-9     CKD Base -      P3      10017

dsccli> rmckdvol -quiet BD02-BD03
CMUN02947E rmckdvol: BD02: The Delete logical volume task cannot be initiated because the
Allow Host Pre-check Control Switch is set to true and the volume that you have specified
is online to a host.
CMUC00024I rmckdvol: CKD volume BD03 successfully deleted.

dsccli> lsckdvol BD02-BD03
Name          ID  accstate  datastate  configstate  deviceMTM  voltype  orgbvols  extpool  cap (cyl)
=====
CKD_EAV1_BD02 BD02 Online    Normal    Normal     3390-9     CKD Base -      P3      10017

dsccli> rmckdvol -force BD02
CMUC00023W rmckdvol: The alias volumes associated with a CKD base volume are automatically
deleted before deletion of the CKD base volume. Are you sure you want to delete CKD volume
BD02?
[Y/N]: y
CMUC00024I rmckdvol: CKD volume BD02 successfully deleted.

dsccli> lsckdvol BD02-BD03
CMUC00234I lsckdvol: No CKD Volume found.
```

Reinitializing online space-efficient CKD volumes

You can use the **initckdvol** command to reinitialize online space-efficient CKD volumes.

The command includes options to prevent the accidental reinitialization of volumes that are in use. A CKD volume is considered to be in use if it is participating in a CS relationship or if the IBM Z system path mask indicates that the volume is in a grouped state or online to any host system. All data is lost when this command is used.

Volume reinitialization is controlled by the **-action releasespace** and **-force** parameters in the following manner:

- ▶ The **releasespace** parameter is used to free all extents and tracks that are associated with the specified volume so they can be reused by other space efficient volumes.
- ▶ The **-force** parameter is required in order for the **initckdvol** command to reinitialize the specified volume.
- ▶ When both of these parameters are used, the user is prompted with a Y/N response in order for the command to proceed with the reinitialization process, as shown in Example 10-68.

Example 10-68 An initckdvol command example

```
dsccli> initckdvol -action releasespace -force BD06
CMUC00338W initckdvol: Are you sure that you want to free all extents and lose the
data associated with CKD volume BD06? [Y/N]:y
CMUC00340I initckdvol: BD06: The command releasespace has completed successfully.
```

10.4.5 Resource groups

The resource group feature is designed for multitenancy environments. The resources are volumes, LCUs, and LSSs, and they are used for access control for CS functions only.

For more information about resource groups, see *IBM System Storage DS8000 Copy Services Scope Management and Resource Groups*, REDP-4758.

10.4.6 IBM Easy Tier

Easy Tier is designed to automate data placement throughout the storage pool. It enables the system, without disrupting applications, to relocate data (at the extent level) across up to three storage tiers. The process is fully automated. Easy Tier also automatically rebalances extents among ranks within the same tier, removing the workload skew between ranks, even within homogeneous and single-tier extent pools.

Easy Tier Heat Map Transfer (HMT) allows the transfer of Easy Tier heat maps from primary to auxiliary storage sites.

For more information about Easy Tier, see the following publications:

- ▶ *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667
- ▶ *IBM DS8870 Easy Tier Heat Map Transfer*, REDP-5015

10.5 Metrics with DS CLI

This section describes several command examples from the DS CLI that analyze the performance metrics from different levels in a storage unit. The DS GUI also provides new capabilities for performance monitoring, as described in 9.14, “Performance monitoring” on page 322.

Important: The `help` command shows specific information about each of the metrics.

Performance metrics: All performance metrics are an accumulation starting from the most recent counter-wrap or counter-reset. The performance counters are reset on the following occurrences:

- ▶ When the storage unit is turned on.
- ▶ When a server fails and the failover and fallback sequence is run.

Example 10-69 shows an example of the `showfbvol` command. This command displays the detailed properties for an individual volume and includes a `-metrics` parameter that returns the performance counter-values for a specific volume ID.

Example 10-69 Metrics for a specific fixed-block volume

```
dscli> showfbvol -metrics 000A
ID                               000A
Date                             10/10/2019 14:19:03 CEST
normrdqrts                       7129645
normrdhits                       4725660
normwritereq                     25533232
normwritehits                   25533232
seqreadreqs                     11746083
seqreadhits                     11578653
seqwritereq                     3994518
seqwritehits                    3994518
cachfwrreqs                     0
cachfwrhits                     0
cachfwreqs                      0
cachfwhits                      0
inbcachload                     0
bypasscach                      0
DASDtrans                       4019246
seqDASDtrans                    5825774
cachetrans                      14933989
NVSspadel                       0
normwriteops                    0
seqwriteops                    0
reccachemis                    1314738
qwriteprots                    0
CKDirtrak                      0
CKDirtrkhits                   0
cachspdelay                    0
timelowifact                   0
thread                         9848943
phwrite                         820103
phbyteread                     4269617
phbytewrite                    7453693
```

```

recmoreads          2104182
sfiletrkreads      0
contamwrts         0
PPRCtrks           0
NVSpallo           29527749
timephread         508529
timephwrite        169907
byteread           4254384
bytewrit           7690334
timbered           336121
typewrite          1418265
zHPFRead           -
zHPFWrite          -
zHPFPrefetchReq   0
zHPFPrefetchHit   0
GMCollisionsSidefileCount 0
GMCollisionsSendSyncCount 0
dscli>

```

Example 10-70 show an example of the **showckdvol** command. This command displays the detailed properties for an individual volume and includes a **-metrics** parameter that returns the performance counter-values for a specific volume ID.

Example 10-70 Metrics for a specific CKD volume

```

dscli> showckdvol -metrics 7b3d
ID          7B3D
normrdqrts  9
normrdhits  9
normwritereq 0
normwritehits 0
seqreadreqs 0
seqreadhits 0
seqwritereq 0
seqwritehits 0
cachfwrreqs 0
cachfwrhits 0
cachfwreqs 0
cachfwhits 0
inbcachload 0
bypasscach 0
DASDtrans   201
seqDASDtrans 0
cachetrans  1
NVSpade1    0
normwriteops 0
seqwriteops 0
reccachemis 0
qwriteprots 0
CKDirtrkac  9
CKDirtrkhits 9
cachspdelay 0
timelowifact 0
phread      201
phwrite     1
phbyteread  49

```

```

phbytwrite          0
recmoreads          0
sfiletrkreads      0
contamwrts         0
PPRCtrks           0
NVSspallo          0
timephread         90
timephwrite        0
byteread           0
bytewrit           0
timeread           0
timewrite          0
zHPFRead           0
zHPFWrite          0
zHPFPrefetchReq    0
zHPFPrefetchHit    0
GMCollisionsSidefileCount 0
GMCollisionsSendSyncCount 0

```

Example 10-71 shows an example of the output of the **showrank** command. This command generates two types of reports. One report displays the detailed properties of a specified rank, and the other report displays the performance metrics of a specified rank by using the **-metrics** parameter.

Example 10-71 Metrics for a specific rank

```

dscli> showrank -metrics R0
ID          R0
Date       10/10/2019 14:32:46 CEST
byteread   8175947
bytewrit   29508105
Reads      17783022
Writes     17946281
timeread   1105024
timewrite  19434052
dataencrypted yes
timeload   8206273
dscli>

```

Example 10-72 shows an example of the **showioport** command. This command shows the properties of a specified I/O port and the performance metrics by using the **-metrics** parameter. Monitoring the I/O ports is one of the most important tasks of the system administrator. The I/O port is where the HBAs, SAN, and DS8900F exchange information. If one of these components has problems because of hardware or configuration issues, all of the other components are affected.

Example 10-72 Metrics for a specific 32 Gbps I/O port

```

dscli> showioport -metrics I0032
ID          I0032
Date       10/12/2019 09:33:38 CEST
byteread (FICON/ESCON) 0
bytewrit (FICON/ESCON) 0
Reads (FICON/ESCON)    0
Writes (FICON/ESCON)   0
timeread (FICON/ESCON) 0

```

timewrite (FICON/ESCON)	0
CmdRetries (FICON)	0
TransferReady (FICON)	0
Logins (FC)	6
SecCapableLogins (FC)	0
AuthLogins (FC)	0
EncryptedLogins (FC)	0
bytewrit (PPRC)	0
byteread (PPRC)	0
Writes (PPRC)	0
Reads (PPRC)	0
timewrite (PPRC)	0
timeread (PPRC)	0
byteread (SCSI)	19858294
bytewrit (SCSI)	8510622
Reads (SCSI)	323762119
Writes (SCSI)	3058836
timeread (SCSI)	585195
timewrite (SCSI)	1750344
LinkFailErr (FC)	0
LossSyncErr (FC)	0
LossSigErr (FC)	0
PrimSeqErr (FC)	0
InvTxWordErr (FC)	0
CRCErr (FC)	0
LRSent (FC)	0
LRRec (FC)	0
IllegalFrame (FC)	0
OutOrdData (FC)	0
OutOrdACK (FC)	0
DupFrame (FC)	0
InvRelOffset (FC)	0
SeqTimeout (FC)	0
BitErrRate (FC)	0
RcvBufZero (FC)	0
SndBufZero (FC)	0
RetQFullBusy (FC)	0
ExchOverrun (FC)	0
ExchCntHigh (FC)	0
ExchRemAbort (FC)	0
CurrentSpeed (FC)	8 Gbps
%UtilizeCPU (FC)	2 Dedicated
TxPower(RDP)	-1.7 dBm(681.2 uW)
RxPower(RDP)	-3.5 dBm(446.8 uW)
TransceiverTemp(RDP)	40 C
SupplyVolt(RDP)	3367.0 mV
TxBiasCurrent(RDP)	6.55 mA
ConnectorType(RDP)	SFP+
TxType(RDP)	Laser-SW
FECStatus(RDP)	Inactive
UncorrectedBlks(RDP)	-
CorrectedBlks(RDP)	-

The output of the **showioport** command includes several metric counters. For example, the `%UtilizeCPU` metric for the CPU utilization of the HBA and the `CurrentSpeed` that the port uses might be useful information.

Example 10-72 on page 393 shows the many important metrics that are returned by the command. It provides the performance counters of the port and the FC link error counters. The FC link error counters are used to determine the health of the overall communication.

The Security Login Counts metric includes the following metrics:

- ▶ **Logins:** The number of remote hosts or peer `N_ports` logged in to the specified I/O port.
- ▶ **SecCapableLogins:** The number of logins that are capable of FC security protocols. If the port is security-enabled, the remote ports are expected to perform authentication.
- ▶ **AuthLogins:** The number of security capable logins that negotiated link authentication.
- ▶ **EncryptedLogins:** The number of remote hosts or peer `N_ports` that are logged in to the specified port. These ports successfully authenticated and negotiated an encryption algorithm for the IBM Fibre Channel Endpoint Security function.

The following groups of errors point to specific problem areas:

- ▶ Any nonzero figure in the counters `LinkFailErr`, `LossSyncErr`, `LossSigErr`, and `PrimSeqErr` indicates that unstable HMCs might be attached to the SAN. These unstable HBAs log in to and log out of the SAN, creating name server congestion and performance degradation.
- ▶ If the `InvTxWordErr` counter increases by more than 100 each day, the port is receiving light from a source that is not a small form-factor pluggable (SFP). The cable that is connected to the port is not covered at the end or the I/O port is not covered by a cap.
- ▶ The `CRCErr` counter shows the errors that occur between the last sending SFP in the SAN and the receiving port of the DS8900F. These errors do not appear in any other place in the data center. To resolve this issue, replace the cable that is connected to the port or the SFP in the SAN.
- ▶ The link reset counters `LRSent` and `LRRec` also suggest hardware defects in the SAN. These errors must be investigated.
- ▶ The counters `IllegalFrame`, `OutOrdData`, `OutOrdACK`, `DupFrame`, `InvRelOffset`, `SeqTimeout`, and `BitErrRate` point to congestion in the SAN. These counters can be influenced only by configuration changes in the SAN.
- ▶ The 16 and 32 Gbps host adapters implement the T11 Read Diagnostic Parameters (RDP) standard and provide extra details, such as optical signal strength, error counters, and other information that is crucial to determining the quality of the link. For example, the cable - connector path (including the cleanliness of the optical connectors) is diagnosed by calculating the `RxPower (RDP)/TxPower (RDP)` ratio. Receivers rarely fail, and the receiver sensitivity does not change. Therefore, if the receiver optical power is too low for good signal reception and the calculated `RxPower (RDP)/TxPower (RDP)` ratio is too low, clean the connector. If this RX/TX ratio continues to be low, the cable might be broken.

- ▶ The following **ioport** command options were added to improve the RDP diagnostic capabilities:
 - **showioport -rdp Ixxxx**
 - Shows detailed RDP information for both the local port and the attached switched port or directly connected port.
 - The information that is shown is the last time that it was obtained from the SFP for the local port or from the attached port that uses RDP.
 - DS8900F 16 Gbps and 32 Gbps FC host adapters read the SFPs each hour and send an RDP command to the attached port every 4 hours.
 - **setioport -update rdp Ixxxx**
 - Causes the port to read the local SFP and also send an RDP command to the attached port.
 - Attached port information currency depends on the implementation.

Example 10-73 provides an example of the RDP status.

Example 10-73 Full output for the showioport -rdp Ixxxx command for a specific I/O port

```

dscli> showioport -rdp I0032
ID I0032
WWPN 5005076309039462
Attached WWPN 200D00051EF0EC72
Physical Type FC-FS-3
Link Failure Error 0
Loss of sync Error 0
Loss of Signal Error 0
Primitive Sequence Error 0
Invalid Transmission Word Error 0
CRC Error 0
FEC Status Inactive
Uncorrected Blocks -
Corrected Blocks -
Port Speed Capabilities 8GFC 16GFC 32GFC
Port Operating Speed 8GFC
Advertised B-B Credit 90
Attached Port B-B Credit 8
Nominal RTT Link Latency Unknown
Connector Type SFP+
Tx Type Short Wave Laser
Transceiver Temperature 39.9 C [Operating Range -128 - +128 C]
Tx Bias Current 6.5 mAmps [Operating Range 0 - 131 mAmps]
Transceiver Supply Voltage 3364.4 mV [Operating Range 0 - 3600 mVolts]
Rx Power 448.8 uW(-3.5 dBm) [Operating Range 0 - 6550 uW]
Tx Power 681.3 uW(-1.7 dBm) [Operating Range 0 - 6550 uW]
Last SFP Read time 10/12/2019 09:48:04 CEST
=====SFP Parameters Alarm Levels=====
Element High Warning Low Warning High Alarm Low Alarm
=====
Transceiver Temperature 0 0 0 0
Tx Bias Current 0 0 0 0
Transceiver Supply Voltage 0 0 0 0
Tx Power 0 0 0 0
Rx Power 0 0 0 0
=====Attached Port=====
ID N/A
WWPN 200D00051EF0EC72
Attached WWPN 5005076309039462
Physical Type FC-FS-3
Link Failure Error 0
Loss of sync Error 3

```



```

Loss of Signal Error          2
Primitive Sequence Error      0
Invalid Transmission Word Error 0
CRC Error                     0
FEC Status                    Inactive
Uncorrected Blocks            -
Corrected Blocks              -
Port Speed Capabilities       1GFC 2GFC 4GFC 8GFC
Port Operating Speed          8GFC
Advertised B-B Credit         0
Attached Port B-B Credit      0
Nominal RTT Link Latency      Unknown
Connector Type                SFP+
Tx Type                       Short Wave Laser
Transceiver Temperature       39.0 C [Operating Range -128 - +128 C]
Tx Bias Current               9.0 mAmps [Operating Range 0 - 131 mAmps]
Transceiver Supply Voltage    3281.1 mV [Operating Range 0 - 3600 mVolts]
Rx Power                      690.2 uW(-1.6 dBm) [Operating Range 0 - 6550 uW]
Tx Power                      479.8 uW(-3.2 dBm) [Operating Range 0 - 6550 uW]
Last SFP Read time           10/12/2019 08:40:30 CEST
=====SFP Parameters Alarm Levels=====
Element                       High Warning Low Warning High Alarm Low Alarm
=====
Transceiver Temperature       0           0           0           0
Tx Bias Current               0           0           0           0
Transceiver Supply Voltage    0           0           0           0
Tx Power                      0           0           0           0
Rx Power                      0           0           0           0
dscli>

```

10.5.1 Offload performance data and other parameters

There is an option to offload performance data by using the DS CLI. The command is **offloadfile -perfdata**. It contains performance data that includes one week of data at 1-minute intervals for the system, arrays, pools, and I/O port. Here is the list of parameters for the **offloadfile** command:

- auditlog** Exports one file that contains the audit log for the Management Console (MC) server.
- sdocert** Exports the IBM Certified Secure Data Overwrite (SDO) certificate for each storage facility.
- etdata** Exports two files that contain the Easy Tier summary data.
- config** Offloads two files: One file contains the advanced settings, and another file contains the Install Corrective Service (ICS) settings. Only one file set parameter must be specified.
- pepackage** Generates a PE package that is sent to IBM, but only if you do not specify **-showstatus**. Specify only one file set.
- statesave** Generates a file that is sent to IBM, but only if you do not specify **-showstatus** or **-list**. Specify only one file set.
- odd** Generates a file that is sent to IBM, but only if you do not specify **-showstatus** or **-list**. Specify only one file set.
- perfdata** Downloads the performance summary comma-separated values (CSV) file.
- sysdata** Downloads the system summary file. Specify only one file set.

- etdataCSV** Downloads a compressed file that contains one or more Easy Tier CSV files and an Excel tool that is used for the CSV files. (The CSV files are Microsoft Excel files.) The compressed file is downloaded to a directory that you specify. Specify only one file set.
- fcdata** Exports the Fibre Channel Connectivity report.

Restriction: The **offloadfile** command is not supported by the embedded DS CLI.

The result of the command in Example 10-74 is a *.csv file* with detailed information. For more information, see Figure 10-9.

Example 10-74 The offloadfile command that is issued from DS CLI

```
dscli> offloadfile -perfddata c:\temp
CMUC00428I offloadfile: The perfddata file has been offloaded to
c:\temp\FORMATTED_PERF_SAMPLES_75HAL91.csv.
dscli>
```

```
Storage System
Time","Interval (seconds)","Name/ID","IOPS (IOPS) Read","IOPS (IOPS) Write","IOPS (IOPS) Total","Reads (%) of Total IOPS","L
"2019-10-13 14:24:53.189","60.2","IBM.2107-75HAL91","3.73","0.03","3.76","0.99","0.05","11.51","0.14","8.11","418.71","11.39
"2019-10-13 14:25:53.351","60.2","IBM.2107-75HAL91","3.53","0.02","3.56","0.99","0.05","11.62","0.12","7.82","449.61","10.7"
"2019-10-13 14:26:53.521","60.2","IBM.2107-75HAL91","4.1","0.04","4.14","0.99","0.05","10.71","0.14","7.53","409.37","11.
"2019-10-13 14:27:53.684","60.2","IBM.2107-75HAL91","3.38","0.02","3.4","0.99","0.05","12.57","0.12","7.69","485.29","10.
"2019-10-13 14:28:53.849","60.2","IBM.2107-75HAL91","3.86","0.03","3.89","0.99","0.05","10.7","0.14","8.19","394.39","11.4",
"2019-10-13 14:29:54.012","60.2","IBM.2107-75HAL91","3.98","0.03","4.01","0.99","0.05","9.72","0.12","7.38","397.91","10.17"
"2019-10-13 14:30:54.193","60.2","IBM.2107-75HAL91","3.47","0.03","3.5","0.99","0.05","12.67","0.14","8.06","466.55","11.
"2019-10-13 14:31:54.354","60.2","IBM.2107-75HAL91","3.84","0.03","3.87","0.99","0.05","10.62","0.13","8.02","401.92","11.19
"2019-10-13 14:32:54.537","60.2","IBM.2107-75HAL91","3.44","0.02","3.47","0.99","0.05","12.53","0.13","7.65","471.05","10.88
"2019-10-13 14:33:54.695","60.2","IBM.2107-75HAL91","4.01","0.03","4.04","0.99","0.05","10.48","0.13","7.82","400.77","10.99
"2019-10-13 14:34:54.867","60.2","IBM.2107-75HAL91","3.78","0.03","3.81","0.99","0.05","10.19","0.12","7.7","405.23","10.
"2019-10-13 14:35:55.027","60.2","IBM.2107-75HAL91","3.52","0.03","3.55","0.99","0.05","12.54","0.15","8.03","455.84","11.55
"2019-10-13 14:36:55.188","60.2","IBM.2107-75HAL91","3.94","0.03","3.97","0.99","0.05","9.7","0.12","7.77","390.58","10.5
"2019-10-13 14:37:55.377","60.2","IBM.2107-75HAL91","3.41","0.02","3.44","0.99","0.05","13.03","0.14","7.69","478.48","11.05
"2019-10-13 14:38:55.537","60.2","IBM.2107-75HAL91","3.99","0.03","4.02","0.99","0.05","10.07","0.12","7.9","390.8","10.8
"2019-10-13 14:39:55.698","60.2","IBM.2107-75HAL91","3.78","0.03","3.81","0.99","0.05","12.27","0.14","7.46","421.26","10.59
"2019-10-13 14:40:55.869","60.2","IBM.2107-75HAL91","3.59","0.03","3.62","0.99","0.05","13.67","0.16","8.16","449.54","11.62
"2019-10-13 14:41:56.05","60.2","IBM.2107-75HAL91","3.98","0.03","4.01","0.99","0.05","10.22","0.12","7.64","397.21","10.54"
```

Figure 10-9 Snapshot from the *-perfddata csv file*

10.6 Private network security commands

DS CLI commands are available that can be used to manage network security on the DS8900F. With the introduction of support for National Institute of Standards and Technology (NIST) 800-131a compliance, new commands were introduced to enable compliance support. Network security includes Transport Layer Security (TLS) to protect application access.

The following commands are available to manage TLS security settings:

- ▶ **manageaccess**

The **manageaccess** command is used to manage the security protocol access settings of an HMC for all communications to and from the DS8900F. It can also control port 1750 access to the network interface server for pre-Gen 2 certificate access.

It is primarily used to manage server access in the HMC, which includes DS GUI, Web User Interface (WUI) and network interface servers.

Each of these accesses can be set to an access level of either Legacy or 800131a. When the access is set to the 800131a level, it is NIST SP 800-131a-compliant.

► **showaccess**

This command displays the current setting for each access that is managed with the **manageaccess** command. It also displays the remote service access settings that are provided with the **lsaccess** command.

The following security commands are available to manage remote service access settings:

► **chaccess**

Use the **chaccess** command to change the following settings of an HMC:

- Enable and disable the command-line shell access to the HMC.
- Enable and disable the WUI access on the HMC.
- Enable and disable Assist On-site (AOS) or IBM Remote Support Center (RSC) access to the HMC.

Important:

- This command affects service access only and does not change access to the system by using the DS CLI or DS Storage Manager.
- Only users with administrator authority can access this command.

► **lsaccess**

The **lsaccess** command displays the access settings of an HMC. If you add the **-l** parameter, the command also displays the AOS or RSC status. If AOS or RSC is active, an AOS or RSC connection shows as enabled. An AOS or RSC connection is used only for remote support purposes. For more information, see Example 10-75.

Example 10-75 The lsaccess command and example output

```
dscli> lsaccess -l
hmc  cmdline  wui      modem    cim      aos      rsc      vpn
=====
hmc1 enabled  enabled  disabled disabled enabled  enabled  disabled
hmc2 enabled  enabled  disabled disabled enabled  enabled  disabled
```

The following commands enable the TLS protocol for secure syslog traffic. TLS must be enabled before configuring all syslog servers. If you specify TLS, all syslog servers configurations use the same protocol and certificates.

► **mksyslogserver**

Example 10-76 shows the new DS CLI command **mksyslogserver**, which configures syslogserver as TLS-enabled. The certificate authority (CA) certificate, HMC Certificate, and HMC private key locations are required when configuring the first syslogserver.

Example 10-76 Enabling TLS for secure syslog traffic

```
dscli> mksyslogserver -addr 9.100.10.7 -protocol tls -cacert
/Users/heping/ut/syslog/ca.pem -hmc-cert /Users/heping/ut/syslog/cert.pem -key
/Users/heping/ut/syslog/key.pem test
CMUC00508I mksyslogserver: The syslog server machine.example.net has been
created.
```

► **lssyslogserver -l**

The **lssyslogserver -l** displays the list of all syslog servers and their attributes, as shown in Example 10-77.

Example 10-77 The lssyslogserver -l command and a sample output

```
dsccli> lssyslogserver -l
name IP address  port state  access protocol type                HMC
=====
daisy 9.10.100.127 514 active online  tls          audit,message,event 1
daisy 9.10.100.127 514 active online  tls          audit,message,event 2
```

Important: For more information about security issues and overall security management to implement NIST 800-131a compliance, see *IBM DS8870 and NIST SP 800-131a Compliance*, REDP-5069.

The following DS CLI commands specify a custom certificate for communication between the encryption key servers (typically IBM Security Key Lifecycle Manager and the storage system):

► **managekeygrp -action -importcert**

Specifies the customer-generated certificate to import in Public-Key Cryptography Standard (PKCS) #12 format, as shown in Example 10-78.

Example 10-78 The managekeygrp command for importing a custom certificate

```
dsccli> managekeygrp -action importcert -loc /home/hscroot/rashmic/da6.p12 -pw blah 1
CMUC00489I managekeygrp: The certificate for encryption group 1 has been imported.
```

► **managekeygrp -action updatecert**

When specified, the option **-certType CUSTOMER** updates an IBM supplied Gen 1 or Gen 2 certificate to a customer-generated certificate, as shown in Example 10-79. For more information about the encryption and DS8000 certificates, see *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security (DS8000 Release 9.2)*, REDP-4500.

Example 10-79 The managekeygrp command for updating a custom certificate

```
dsccli> managekeygrp -action updatecert -certType CUSTOMER -key data 1
CMUC00472I managekeygrp: The certificate for encryption group 1 has been updated.
```

For more information, see *IBM DS8000 Series Command-Line Interface User's Guide*, SC27-9562.

10.7 Copy Services commands

More DS CLI commands are available, and many of these commands are used for the management of CS, such as the FlashCopy, Metro Mirror (MM), and Global Mirror (GM) commands.

These commands are not described in this chapter. For more information about these commands, see *IBM DS8000 Copy Services: Updated for IBM DS8000 Release 9.1*, SG24-8367.

10.8 Earlier DS CLI commands and scripts

As versions of the DS8900F code evolve, new commands are introduced that are supported by the DS CLI. Additionally, in some cases, older commands are adjusted or changed to support new functions. To ensure continuity for any scripts and programming customers might have created by using these older commands, the DS CLI still supports many of these commands even though they might no longer be documented in the DS CLI Command Reference or available any longer with the DS CLI **help** command, as shown in Example 10-80.

Example 10-80 Command help

```
dscli> help lsddm
Help page for lsddm is not available
dscli>
```

Even though the command may no longer be referenced in the **help** pages of the DS CLI, the command is still supported, as shown in Example 10-81.

Example 10-81 Support for older commands

```
dscli> lsddm
Date/Time: June 27, 2022 9:11:04 PM CEST IBM DSCLI Version: 7.9.30.154 DS: -
ID          DA Pair dkcap (10^9B) dkuse          arsite State
=====
IBM.2107-D04-4N014/R1-P1-D1 11          1920.0 array member  S3      Normal
IBM.2107-D04-4N014/R1-P1-D2 11          1920.0 array member  S3      Normal
IBM.2107-D04-4N014/R1-P1-D3 11          1920.0 array member  S3      Normal
IBM.2107-D04-4N014/R1-P1-D4 11          1920.0 array member  S3      Normal
```

Some new commands and their older equivalents are shown in Table 10-1.

Table 10-1 DS CLI new and old commands

New command	Output	Old command
lsdrive	Lists drives.	lsddm
lsstgenclosure	Lists storage enclosures.	lsstgenc1
lshacard	Lists host adapters.	lshba
lsdacard	Lists device adapter (DA).	lsda
lspnode	Lists processor nodes.	N/A
lshmc	Lists HMCs.	N/A
lsioenclosure	Lists I/O enclosures.	N/A
lspciocard	Lists PCIe cards.	N/A

10.9 For more information

For more information about the topics that were described in this chapter, see the following resources:

- ▶ For more information about the CS configuration, see *IBM DS8000 Series Command-Line Interface User's Guide*, SC27-9562.
- ▶ For more information about DS CLI commands that relate to disk encryption, see *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security (DS8000 Release 9.2)*, REDP-4500.
- ▶ For more information about thin-provisioning, see *IBM DS8880 Thin Provisioning (Updated for Release 8.5)*, REDP-5343.
- ▶ For more information about DS CLI commands that relate to Easy Tier, see *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667.



Part 4

Maintenance and upgrades

This part of the book provides useful information about maintenance and upgrades.

This part contains the following chapters:

- ▶ Chapter 11, “Licensed Machine Code” on page 405
- ▶ Chapter 12, “Monitoring and support” on page 423.



Licensed Machine Code

This chapter describes considerations that relate to the planning and installation of new Licensed Machine Code (LMC) bundles on the IBM DS8900F. The overall process for the DS8900F is the same as for previous models. However, several enhancements to IBM Power firmware updates are described.

This chapter covers the following topics:

- ▶ How new Licensed Internal Code is released
- ▶ Bundle installation
- ▶ Code updates
- ▶ Host adapter firmware updates
- ▶ Loading the code bundle
- ▶ Fast path concurrent code load
- ▶ Postinstallation activities
- ▶ Summary

11.1 How new Licensed Internal Code is released

Most of the hardware components within the DS8900F system can be updated with new firmware when it is available:

- ▶ Device adapters (DAs)
- ▶ Host adapters
- ▶ Power subsystems:
 - Intelligent Power Distribution Units (iPDUs)
 - Rack power control
 - Fibre Channel Interface Cards (FCICs)

In addition, the Licensed Internal Code (LIC) and internal operating system (OS) that run on the Hardware Management Consoles (HMCs) and each central processor complex (CPC) can be updated. As IBM continues to develop the DS8900F, new features are released through new LIC levels.

When IBM releases new LIC for the DS8900F, it is released in the form of a bundle. The term *bundle* is used because a new code release can include updates for various DS8900F components. These updates are tested together, and then the various code packages are bundled together into one unified release. Components within the bundle each include their own revision levels.

For more information about a DS8900F cross-reference table of code bundles, see [DS8900F Code Bundle Information](#).

The cross-reference table shows the levels of code for released bundles. The cross-reference information is updated as new code bundles are released.

In addition to keeping your LIC up to date, make sure to maintain a current version of the Data Storage Command-line Interface (DS CLI).

The DS8900F uses the following naming convention for bundles: PR.MM.FFF.EEEE, where the components are:

- ▶ P: Product (8 = DS8000)
- ▶ R: Release Major (X, 9 = DS8900F)
- ▶ MM: Release Minor (xx)
- ▶ FFF: Fix Level (xxx)
- ▶ EEEE: Fix Level (0 is base, and 1..n is the interim fix build that is later than the base level.)

This naming convention is shown in Example 11-1.

Example 11-1 Bundle level information

For BUNDLE 89.30.68.0 :	
Product	DS8000
Release Major	9
Release Minor	30
Fix Level	68
EFIX level	0

The 9.30 in Example 11-1 stands for the Release 9.3 without a Service Pack.

If DS CLI is used, you can obtain the CLI and LMC code level information by using the `ver` command, as shown in Example 11-2. The `ver` command uses the following optional parameters and displays the versions of the CLI, Storage Manager, and LMC:

- s (optional) The -s parameter displays the version of the CLI program. You cannot use the -s and -l parameters together.
- l (optional) The -l parameter displays the versions of the CLI, Storage Manager, and LMC. You cannot use the -l and -s parameters together.
- cli (optional) Displays the version of the CLI program. Version numbers are in the format *version.release.modification.fixlevel*.
- stgmgr (optional) Displays the version of the Storage Manager.
This ID is not for the GUI (Storage Manager GUI). This ID relates to HMC code bundle information.
- lmc (optional) Displays the version of the LMC.

Example 11-2 Output of the DS CLI ver -l command

```

dscli> ver -l
Date/Time: June 28, 2022 12:10:08 PM CEST IBM DSCLI Version: 7.9.30.154 DS: -
DSCLI          7.9.30.154
StorageManager 0.0.0.0
HMC DSCLI      7.9.30.154
=====Version=====
Storage Image  LMC          Bundle Version
=====
IBM.2107-75HAL91 7.9.30.154 89.30.68.0
  
```

The Bundle version (Release) also can be retrieved from the DS Storage Manager by clicking **Actions** → **Properties** from the Dashboard window, as shown in Figure 11-1.

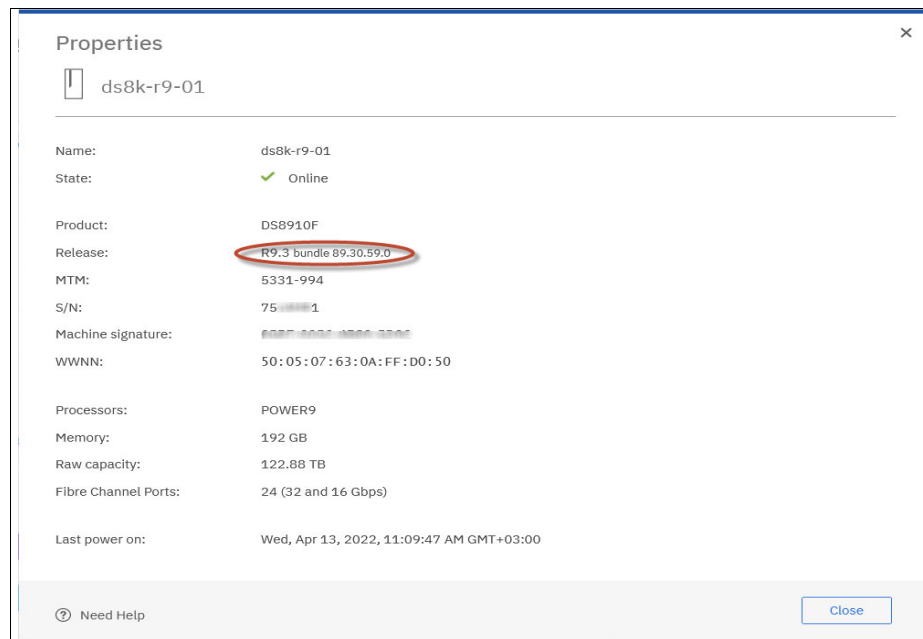


Figure 11-1 Bundle version under DS Storage Manager

11.2 Bundle installation

Important: The LMC is usually provided by and installed by IBM Remote Support Personnel, or by an IBM Systems Service Representative (IBM SSR). With the release of R9.3, the customer may manage the entire process from the DS8000 Storage Manager GUI. When this process is handled by the IBM SSR or IBM Remote Support Personnel, they review the “Prerequisites” section or “Attention Must Read” section in the LIC update instructions, and inform the customer during the planning phase if any prerequisites must be considered.

The bundle package contains the following new levels of updated code:

- ▶ HMC code levels:
 - HMC OS and managed system base
 - DS Storage Manager
 - IBM Copy Services Manager
- ▶ Managed system code levels
- ▶ Interim fix code levels
- ▶ Storage facility image (SFI) code levels
- ▶ Host adapter code levels
- ▶ DA code level
- ▶ I/O enclosure code level
- ▶ Power code levels
- ▶ Rack power control card (RPCC) code level
- ▶ Flash Drive Enclosure Service Module (ESM) interface card code levels
- ▶ Flash Drive enclosure power supply unit (PSU) code levels
- ▶ Flash drive module firmware code level

The code is either updated remotely or locally at the HMC by an IBM SSR. Upgrading the code remotely can be done by IBM through Remote Code Load (RCL) or by the client through the DS8000 Storage Manager GUI. RCL is the default method. If the client wants the code updated by the local IBM SSR onsite, then the Feature Code for remote code exception must be ordered with the system.

Note: When the customer does not opt for Expert Care Premium, Customer Code Load is the default on the DS8900F system.

Other than the actions of acquiring the microcode, the process of distribution and activation is the same.

The Code Distribution and Activation (CDA) software preinstall is the method that is used to run the concurrent code load (CCL) distribution. By using the CDA software preinstall, the IBM SSR performs every non-impacting CCL step for loading code by inserting the physical media into the primary HMC or by running a network acquisition of the code level that is needed. The IBM SSR can also download the bundle to their Notebook and then load it on the HMC by using a service tool, or download the bundle from IBM Fix Central on the HMC for RCL.

After the CDA software preinstallation starts, the following steps occur automatically:

1. The release bundle is downloaded from either the DVD or network to the HMC hard disk drive (HDD).
2. The HMC receives any code update-specific fixes.
3. Code updates are distributed to the logical partition (LPAR) and staged on an alternative base operating system (BOS) repository.
4. Scheduled precheck scans are performed until the distributed code is activated by the user. After 30 days without activation, the code expires and is automatically removed from the alternative BOS.

Anytime after the software preinstallation completes, when the user logs in to the primary HMC, the user is guided automatically to correct any serviceable events that might be open, update the HMC, and activate the previously distributed code on the storage facility. The overall process is also known as CCL.

The code installation process performs the following actions:

1. Updates code on the primary HMC (HMC1).
2. If a dual-HMC configuration is used, the code is acquired from the primary HMC (HMC1) and copied to the secondary HMC (HMC2).
3. Performs updates to the CPC OS and the internal LMC, which are performed individually. To update each CPC, the logical subsystem (LSS) fails over to the alternative CPC. This process updates the firmware that is running on each DA that is owned by that CPC.
4. Performs updates to the host adapters. For DS8900F host adapters, the impact of these updates on each adapter is less than 2.5 seconds and they do not affect connectivity. If an update takes longer, the multipathing software on the host or the control-unit initiated reconfiguration (CUIR) directs I/O to another host adapter. If a host is attached with only a single path or through multipathing to different ports on the same host adapter, connectivity is lost. For more information about host attachment, see 3.3.2, "Host connections" on page 82.
5. New power supply and RPCC firmware are periodically released. New firmware can be loaded into each RPCC directly from the HMC. For more information, see 3.6, "RAS on the power subsystem" on page 97.

During the RPCC firmware update, RPC0 is first placed into Service Mode while updated and verified, then resumed, and then the alternative card RPC1 completes.

6. New firmware for the hypervisor, service processor, system board, and I/O enclosure boards are periodically released. This firmware can be loaded into each device directly from the HMC. Activation of this firmware might require a shutdown and restart of each CPC individually. This process causes each CPC to fail over its LSSs to the alternative CPC. Certain updates do not require this step, or it might occur without processor restarts. For more information, see 3.2, "CPC failover and failback" on page 78.
7. Maintain the latest flash drive firmware because enhancements for efficiency, performance, and reliability are released often. A flash drive firmware update is a concurrent process in the DS8900F series family. Direction to update the firmware by using Install Corrective Service (ICS) can be found in the code installation instructions.

Although the microcode installation process might seem complex, it does not require significant user intervention. IBM Remote Support Personnel normally start the CDA process and then monitors its progress by using the HMC. The customer's experience with upgrading the code of the DS8000 is the same.

Important: For the DS8900F models, DS CLI should be maintained at a current level. Matching the version to the storage facility is not required if the DS CLI version is at a higher level. The higher level can be used to support all other IBM DS8000 models in the environment. For more information, see the release notes or speak to your IBM SSR.

Automatic Code Management

New since Release 9.2 is an optional automatic acquire and preinstallation feature for LIC. This feature automatically downloads the recommended “go to” level of LIC to the HMC, and then it performs an automated preinstallation. This preinstallation reduces the code load time because this level is preinstalled at the beginning of the code load window. You can view the recommended level of LIC by going to [DS8000 Code Recommendation](#).

Important: The default setting for this feature is off, but can be enabled in the Storage Manager GUI. For more information, contact your IBM SSR.

To enable this feature, log in to the Storage Manager GUI, select **Settings** → **System** → **Advanced**, and select **Automatic code management**, as shown in Figure 11-2.

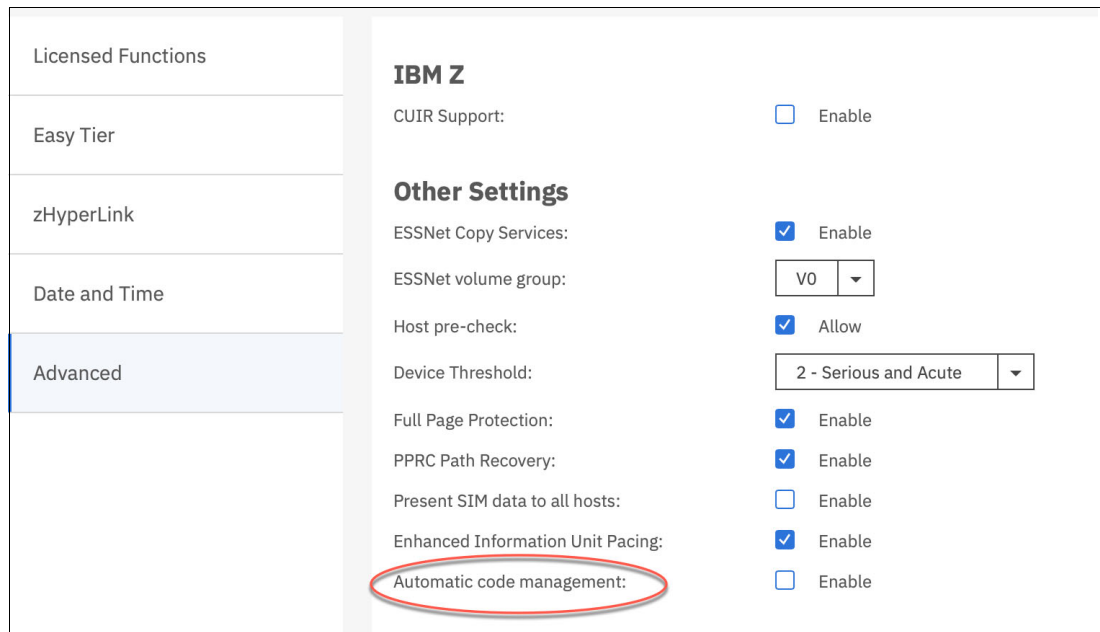


Figure 11-2 Advanced Settings menu: Automatic code management

HMC Code Image Server

When several DS8000 storage systems in the same data center must be updated, acquiring a code bundle from IBM Fix Central individually for each storage system can be a time-consuming process.

To address this situation, a new HMC Code Image Server function was introduced in DS8000 Release 9.0. With the HMC Code Image Server function, a single HMC in the customer data center can acquire code from IBM Fix Central. One HMC sends those images to other HMCs by using the client Ethernet network. The advantage of this approach is that there is no need to download the image from IBM Fix Central multiple times, and the code bundles can be copied locally by using that download.

The HMC Code Image Server function works with bundle images and other updates, such as ICS Images. HMC Recovery Images are performed if they are available on the source HMC.

Figure 11-3 shows the two new menu options in the Updates menu of the service console Web User Interface (WUI).

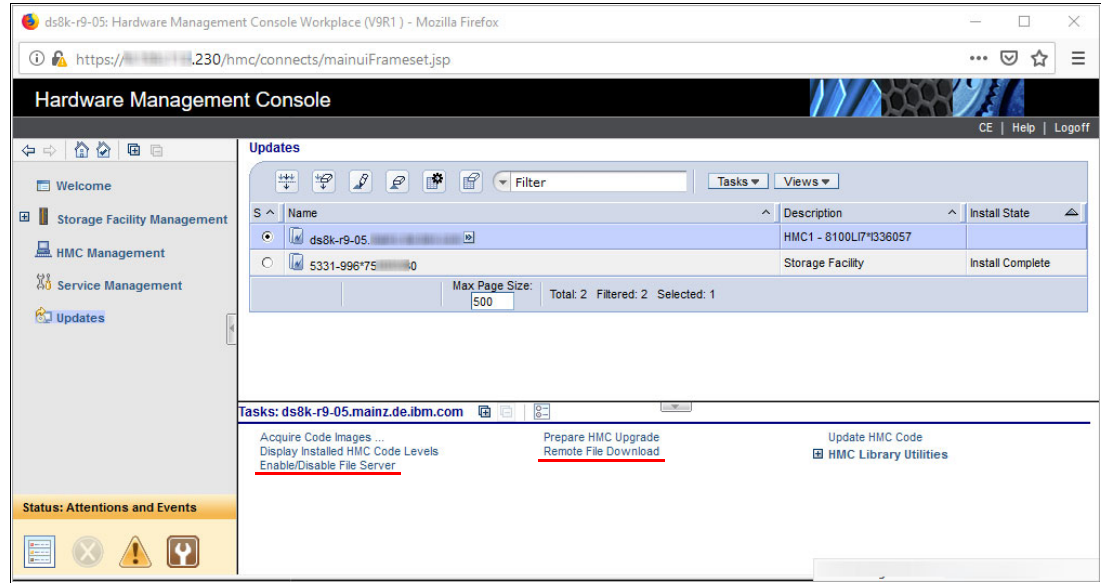


Figure 11-3 Updates menu: HMC Code Image Server

At the site where the code bundle was acquired and downloaded, the HMC Code Image Server function must be enabled. The target site then uses the Remote File Download function to copy the available code bundles to a local repository. All images on the source HMC are copied to the target HMC.

This process copies only the update image files to the local `/extra/BundleImage/` directory on the target HMC. Then, the normal acquisition step still must be performed, and the local directory on the target HMC must be selected, as shown on Figure 11-4.

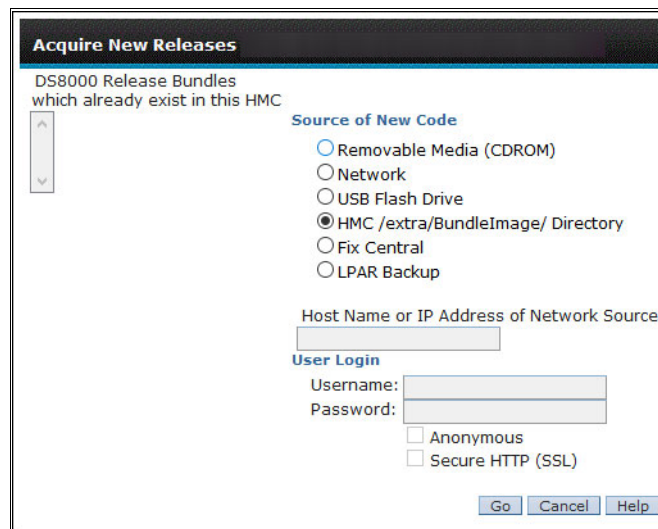


Figure 11-4 Acquiring from the HMC Directory

After the acquisition is complete, the normal code load process proceeds with the CDA software preinstallation.

Note: Bundles that were acquired from physical DVD media cannot be served by the HMC Code Image Server function because they are imported directly to the HMC software library and are not available as single bundle image afterward.

Figure 11-4 on page 411 also shows that it is possible to acquire a bundle from the storage system LPAR. Every IBM Fix Central acquired image is copied to the HMC and the LPARs of the storage system and then imported into the HMC library. Because there is a copy on the LPAR, the partner HMC can now use the LPAR as a source for the acquisition step. This action can be done on both HMCs on the same storage system because only these HMCs have access to the internal network to copy the files from the LPARs. Copying from the LPARs does not require using the HMC Code Image Server function.

11.2.1 Remote Code Load

With RCL, you have an efficient and secure method to update the DS8000 systems microcode in a concurrent way without interrupting business operations.

RCL is supported by DS8870, DS8880, and DS8900F systems.

RCL is a trusted process where an IBM Remote Support engineer securely connects to a DS8000 system, enables the remote acquisition, and performs the distribution and activation of LIC bundles and ICS images.

The RCL process is concurrent, that is, it can be run without interruptions to business operations. This process consists of the following steps, as illustrated in Figure 11-5.

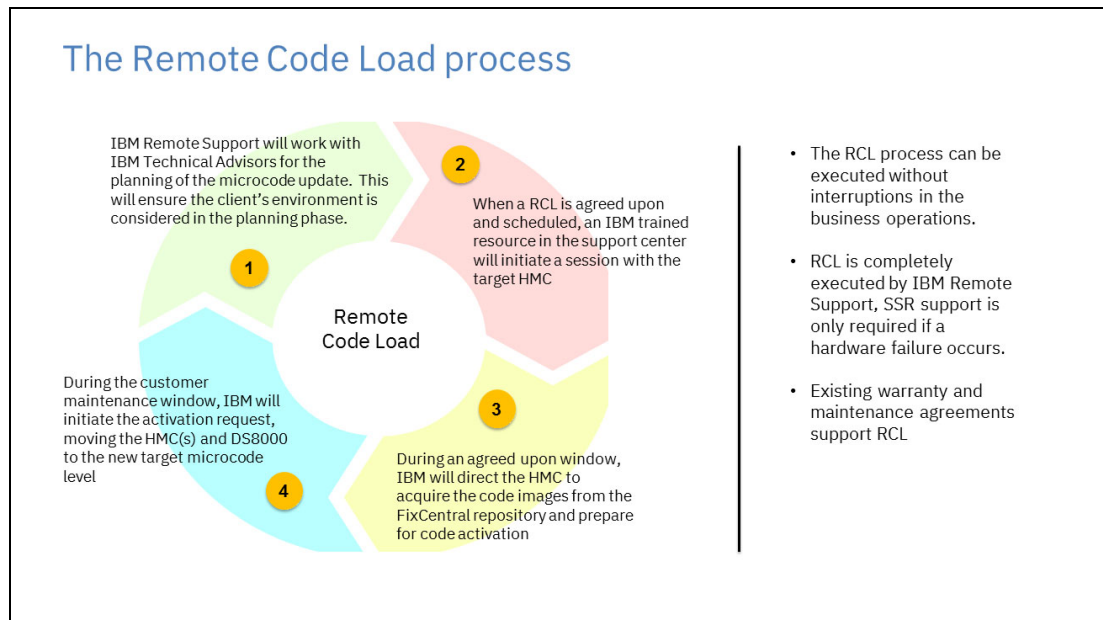


Figure 11-5 Remote Code Load process

1. IBM Remote Support Personnel work with IBM Technical Advisors to plan the microcode update to ensure that the client's environment is in the planning phase.
2. When an RCL is agreed on and scheduled, IBM Remote Support Personnel in the IBM Support Center initiate a session with the target HMC.
3. During the agreed on maintenance window, IBM Remote Support Personnel direct the HMC to acquire the code images from the IBM Fix Central repository and prepare for code activation.
4. During the customer maintenance window, IBM Remote Support Personnel initiate the activation request and update the HMCs and DS8000 to the new target microcode level.

RCL requires port 443 access to the following servers:

- ▶ esupport.ibm.com: 129.42.54.189
- ▶ esupport.ibm.com: 129.42.56.189
- ▶ esupport.ibm.com: 129.42.60.189
- ▶ eccgw01.boulder.ibm.com: 207.25.252.197
- ▶ eccgw02.rochester.ibm.com: 129.42.160.51

Code bundles are pulled to the HMC. They are not pushed.

11.2.2 Customer Code Load

With Release 9.3, the DS8000 introduces a new feature that clients can use to perform an upgrade of the DS8000 storage systems by themselves. This option provides some benefits for clients that prefer having full control of the entire process to plan, schedule, and perform the code upgrade. The process is entirely handled by the customer, and IBM Remote Support Personnel are not required to perform this task.

Note: Customer Code Load runs the same background processes as the RCL.

The process is split into two parts:

1. The user performs the download and acquires the new DS8000 code bundle.
2. The user activates the code bundle.

There is a 30-day countdown between these two parts of the upgrade in which the client can decide when to proceed with the second part. During the 30 days, the system stores the downloaded code bundle, and it can be activated at any point.

A new menu option is available under **Settings** → **Support** → **Update System**, as shown in Figure 11-6.

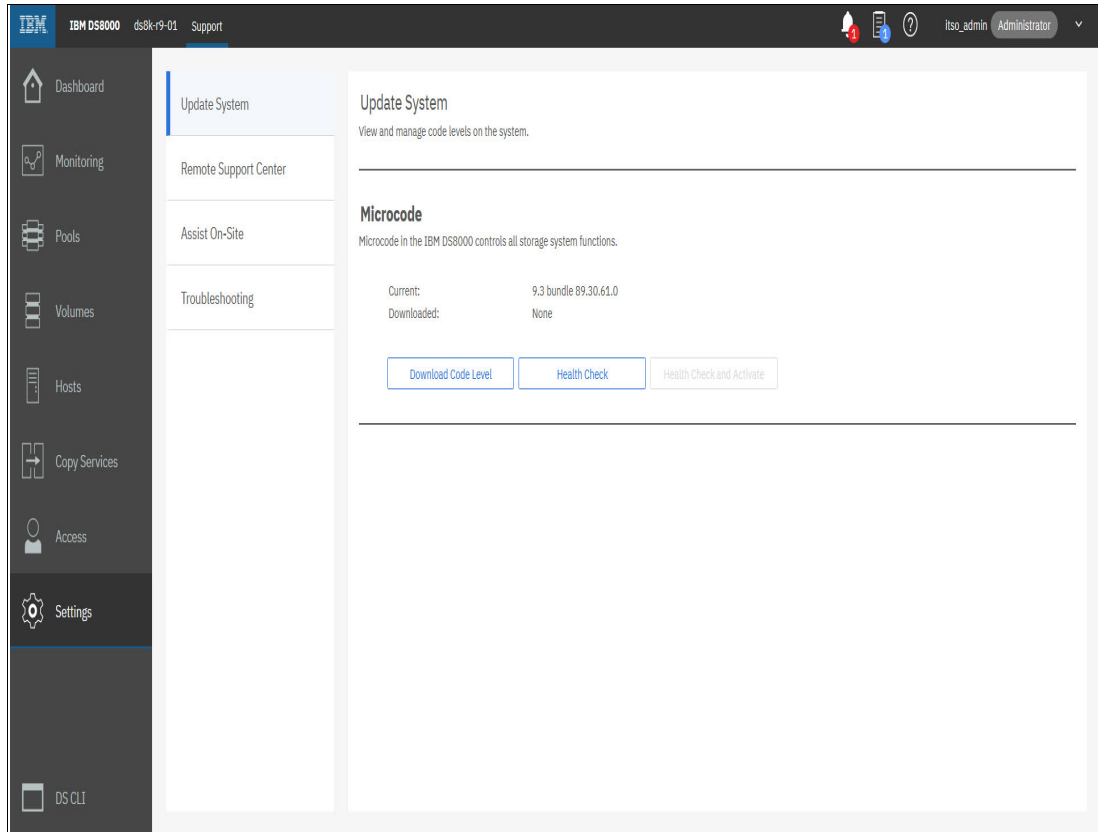


Figure 11-6 Storage Manager GUI: Update System menu

Figure 11-6 shows the Update System menu when there are no code bundles downloaded to the DS8000 storage system. The Health Check and Activate option is enabled after the download of the microcode completes.

Important: The storage system must be in a healthy hardware status to avoid issues during the code upgrade process. For that reason, at any point in time, the client can select the option **Health Check** to confirm whether the system is ready for the upgrade.

Customer Code Load process

To use the Customer Code Load process, complete the following steps:

1. To download the microcode, select **Download Code Level** from the window that is shown in Figure 11-6, which immediately begins a microcode bundle query of IBM Fix Central for a list of codes that the customer can select. This query displays only higher or equal microcode levels compared to the currently installed level on the DS8000 system. After the query completes, a new window opens in the Storage Manager GUI, as shown in Figure 11-7 on page 415.




Figure 11-7 Microcode level query

2. After the code level is selected, the process downloads the new code bundle to the DS8000 Hardware Management Console, and then distributes the separate firmware packages to each internal component in the DS8000 system.
3. The user can monitor the process until completion. After the download completes, click **Close Status**, as shown in Figure 11-8.

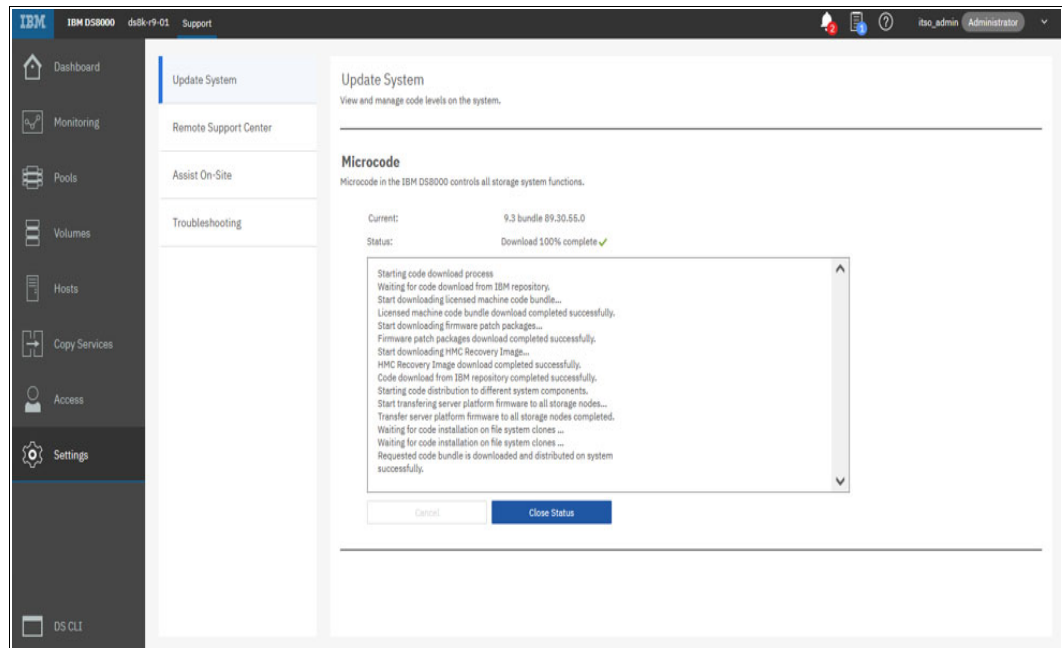


Figure 11-8 Microcode download complete message

4. After closing the window, the option Health Check and Activate is available. Also, you can see and confirm which code level was downloaded, and how many days are left before the code expires, as shown in Figure 11-9.

Note: After completing this step, the Health Check option still is available, but it is not mandatory to select it before proceeding with the code activation. A health check still is performed before and after the activation when the user selects **Health Check and Activate**.

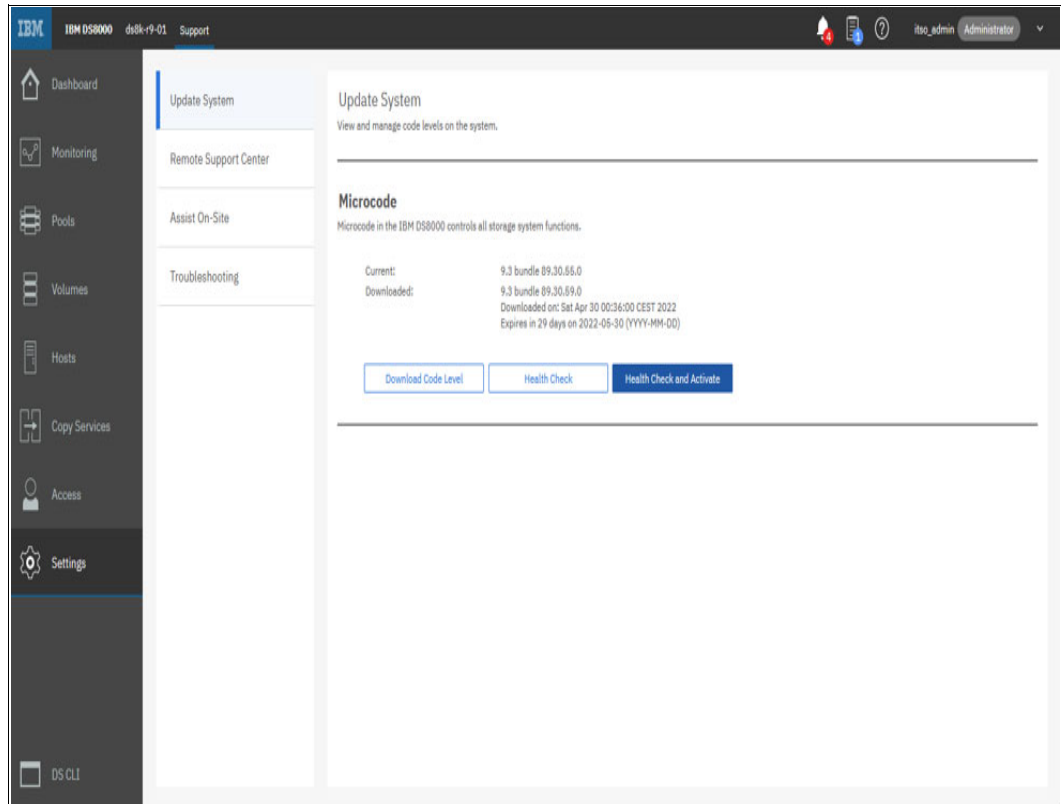


Figure 11-9 Update System: Ready for Activation

5. To activate the code, select **Health Check and Activate**. A new attention message appears in the Storage Manager GUI and notifies you about the actions that are about to be performed, as shown in Figure 11-10 on page 417. To start, click **Yes**.

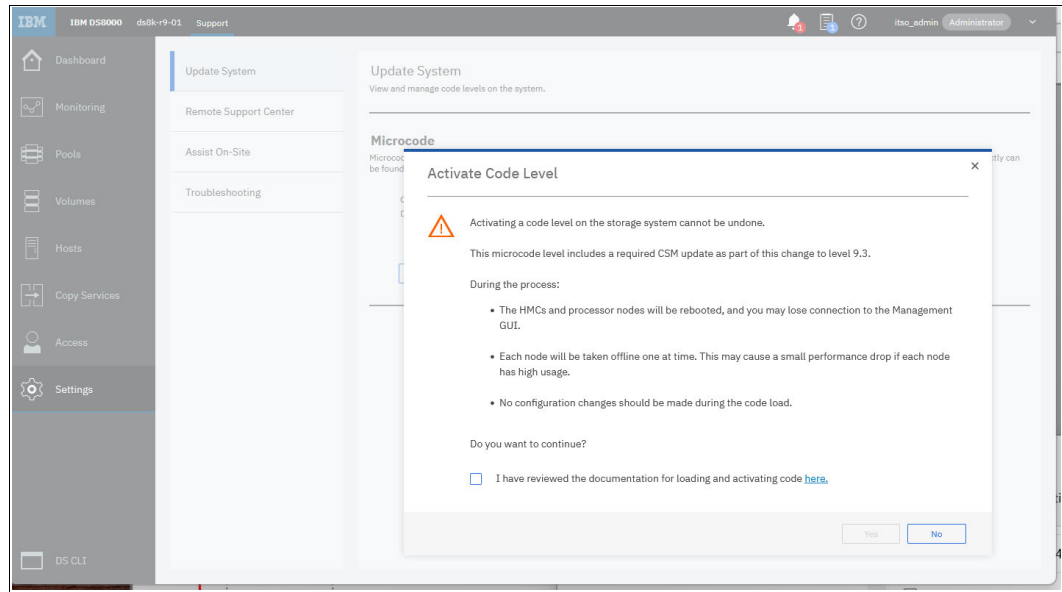


Figure 11-10 Code Activation Message

- The code activation progress can be tracked in the Storage Manager GUI until the end. After it completes, it displays a message confirming that the activation is complete (see Figure 11-11).

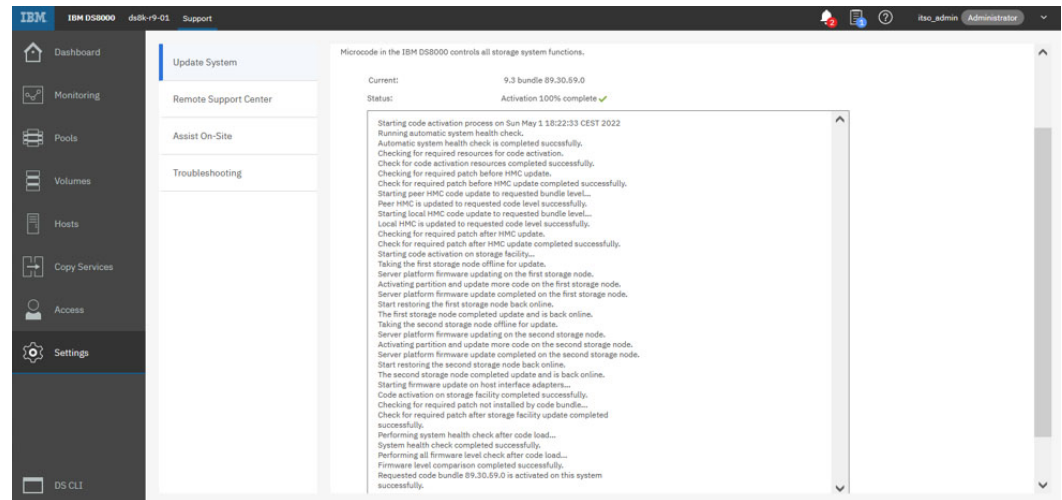


Figure 11-11 Code Activation Complete

If there is any unexpected situation such as hardware failure or code upgrade failure during the process, the customer is notified in the window that is shown in Figure 11-11. There will be a case number in that notification, which can be used as a reference for when the client engages IBM Remote Support.

11.3 Code updates

The LIC running on the HMC is updated as part of a new code bundle installation. The HMC can hold up to six previous versions of code in the Code Library. Each CPC can hold three versions of code (the previous version, the active version, and the next version). Most organizations plan for two code updates each year.

Best practice: Many clients with multiple DS8000 systems follow the update schedule that is detailed in this chapter. In this schedule, the HMC is updated a day or two before the rest of the bundle is applied. If a large gap exists between the present and destination level of bundles, certain DS CLI commands (especially DS CLI commands that relate to IBM Copy Services (CS)) might not be able to be run until the SFI is updated to the same level of the HMC. Your IBM SSR or IBM Technical Advisor or Technical Account Manager can help you in this situation.

Before you update the CPC OS and LIC, a pre-verification test is run to ensure that no conditions exist that prohibit a successful code load. The HMC code update installs the latest version of the pre-verification test. Then, the newest test can be run.

If problems are detected, one or two days are available before the scheduled code installation window date to correct them. This procedure is shown in the following example:

- ▶ Thursday:
 - a. Acquire the new code bundle and send them to the HMCs.
 - b. Update the HMCs to the new code bundle.
 - c. Run the updated pre-verification test.
 - d. Resolve any issues that were identified by the pre-verification test.
- ▶ Saturday:
 - Update (Activate) the SFI code.

The average code load time varies depending on the hardware that is installed, but 2.5 - 4 hours is normal. Always speak with your IBM SSR about proposed code load schedules.

Code recommendations are listed at [DS8000 Code Recommendation](#).

Additionally, check multipathing drivers and storage area network (SAN) switch firmware levels for their current levels at regular intervals.

11.4 Host adapter firmware updates

One of the final steps in the CCL process is updating the host adapters. Normally, every code bundle contains new host adapter firmware. For DS8900F Fibre Channel (FC) adapters, regardless of whether they are used for open systems (FC) attachment or IBM Z system (Fibre Channel connection (IBM FICON)) attachment, the update process is concurrent with the attached hosts. The FC adapters use a technique that is known as *adapter fast-load*, which is the default option for updating. This technique allows the adapters to switch to the new firmware in less than 2 seconds.

This fast update means that single path hosts, hosts that boot from SAN, and hosts that do not have multipathing software do not need to be shut down during the update. They can keep operating during the host adapter update because the update is so fast. Also, no Subsystem Device Driver (SDD) path management is necessary.

Interactive host adapters also can be enabled if you want to control the host path manually. If so, before the host adapters are updated, a notification is sent and a confirmation is needed. You can then take the corresponding host paths offline and switch to other available paths.

11.4.1 Light-on fastload firmware update

Light-on fastload firmware update is a DS8900F feature that allows a host adapter to perform a fastload firmware update without dropping light to hosts and SAN switches. It decreases generated registered state change notification (RSCN) messages in SAN fabrics. It also eliminates warning messages or problems that are generated because of the light-off action of a traditional fastload. In addition, this feature shortens the duration of host I/O interrupts during the host firmware update because the DS8900F host ports stay in login states without extra SAN fabric communication actions, such as a login back to the connected SAN switch.

This function is usually enabled by default. For more information about how to enable this function, contact your IBM SSR.

11.4.2 Remote Mirror and Remote Copy path considerations

No special considerations are required for Remote Mirror and Remote Copy paths that use FC ports. The ability to perform a fast load means that no interruption occurs to the Remote Mirror operations.

11.4.3 Control-unit initiated reconfiguration

CUIR prevents the loss of access to volumes in IBM Z system environments because of incorrect or wrong path handling. This function automates channel path management in IBM Z environments in support of selected DS8900F service actions. CUIR is available for the DS8900F when the DS8900F is operated in the z/OS and z/VM environments. The CUIR function automates channel path vary on and vary off actions to minimize the manual operator intervention during selected DS8900F service actions.

CUIR allows the DS8900F to request that all attached system images set all paths that are required for a particular service action to the offline state. System images with the correct level of software support respond to these requests by varying off the affected paths. The image then notifies the DS8900F subsystem that the paths are offline or that it cannot take the paths offline. CUIR reduces manual operator intervention and the possibility of human error during maintenance actions.

CUIR also reduces the time that is required for the maintenance window. This feature is useful in environments in which many systems are attached to a DS8900F.

11.5 Loading the code bundle

The DS8900F code bundle installation can be performed in different ways. The preferred method is through the IBM RCL process, where IBM Remote Support Personnel download the microcode bundle remotely from IBM Fix Central and then activate it remotely.

Starting with Release 9.3, loading the microcode can now be performed entirely by the client. The microcode is downloaded from [IBM Fix Central](#), and the client can perform all the steps by using the DS8000 Storage Manager GUI. For more information, see 11.2.2, “Customer Code Load” on page 413.

The microcode also can be loaded by an IBM SSR onsite. To review and arrange the required services, contact your IBM SSR or your IBM Technical Account Manager.

11.6 Fast path concurrent code load

DS8900F supports CCL. CCL in DS8900F is basically the same as previous generations of DS8000. CCL is referred to as traditional CCL. DS8000 development always strives to improve the code load function's robustness and reduce activation durations. A faster, more fault-tolerant code load method is known as *fast path concurrent code load* (FPCCL).

FPCCL is automatically set as the preferred code load function, assuming that the requirements of the bundle to be activated satisfy the requirements for FPCCL.

The FPCCL requirements were expanded to include the following features. The delta of the "coming from" level and the "go to" level consists of these elements.

- ▶ SFI code:
 - LPAR code
 - DA
- ▶ High-Performance Flash Enclosure (HPFE):
 - Small Computer System Interface (SCSI) Enclosure Services (SES) processor firmware
 - PSU firmware
- ▶ Host adapter firmware
- ▶ AIX interim fix
- ▶ IBM Power firmware for iPDUs and RPCCs

Important: The code load function reverts to traditional CCL if any additional components, other than the components that are listed previously, are included in the update.

FPCCL includes an *autonomic recovery function*, which means that FPCCL is far more tolerant to temporary non-critical errors that might surface during the activation. During the FPCCL, if an error is posted, the LIC automatically analyzes the error and evaluates whether CCL can continue. If it cannot, the LIC suspends the CCL and calls for service. The DS8900F system can continue with the code update with tolerable errors. The DS8900F FPCCL update is more robust with a much shorter duration. After the code update completes, your IBM SSR works to resolve any of the problems that were generated during the code update at a convenient time, allowing DS8900F clients to schedule the code update in a controlled manner.

During an update, a system is under less redundant conditions because certain components are undergoing a firmware update. With FPCCL, firmware activation time is drastically reduced. Therefore, system redundancy is improved with less exposure to non-redundant durations. In addition, firmware distribution time is also minimized because fewer components are involved in the code update.

The CCL duration of the DS8000 family continues to advance with the introduction of new technology. With the latest DS8900F firmware, the LIC preinstall can be arranged before your code update service window performs the code activation, distribution, and HMC update. The activation times of various components are greatly reduced.

11.7 Postinstallation activities

After a new code bundle is installed, you might need to complete the following tasks:

1. Upgrade the DS CLI of external workstations. For most new release code bundles, a corresponding new release of the DS CLI is available. The LMC version and the DS CLI version are usually identical. Ensure that you upgrade to the new version of the DS CLI to take advantage of any improvements.

A current version of the DS CLI can be downloaded from [IBM Fix Central](#).

2. Verify the connectivity from each DS CLI workstation to the DS8900F.
3. Verify the DS Storage Manager connectivity by using a supported browser.
4. Verify the DS Storage Manager connectivity from IBM Spectrum Control to the DS8900F.
5. Verify the DS Storage Manager connectivity from IBM Copy Services Manager to the DS8900F.
6. Verify the connectivity from the DS8900F to all IBM Security Guardium Key Lifecycle Manager instances, or other servers in use.

11.8 Summary

IBM might release changes to the DS8900F LMC. These changes might include code fixes and feature updates that relate to the DS8900F.

These updates and the information about them are documented in the DS8900F Code cross-reference website. You can find this information for a specific bundle under the Bundle Release Note information section in [DS8000 Code Recommendation](#).



Monitoring and support

This chapter provides information about the Simple Network Management Protocol (SNMP) implementation and messages for the IBM DS8900F storage system.

The chapter also describes the outbound (Call Home and support data offload) and inbound (code download and remote support) communications for the IBM DS8000 family.

This chapter covers the following topics:

- ▶ SNMP implementation on the DS8900F
- ▶ SNMP notifications
- ▶ SNMP configuration
- ▶ Introducing remote support
- ▶ IBM policies for remote support
- ▶ Remote support advantages
- ▶ Remote support and Call Home
- ▶ Remote Support Access (inbound)
- ▶ Call Home and Assist On-Site customer-provided certificates
- ▶ Audit logging
- ▶ Using IBM Storage Insights
- ▶ IBM Call Home Connect Cloud

12.1 SNMP implementation on the DS8900F

SNMP, as used by the DS8900F, is designed so that the DS8900F sends traps only if a notification is necessary. The traps can be sent to a defined IP address.

SNMP alert traps provide information about problems that the storage unit detects. You or the service provider must correct the problems that the traps detect.

The DS8900F does not include an installed SNMP agent that can respond to SNMP polling. The default Community Name parameter is set to `public`.

The management server that is configured to receive the SNMP traps receives all of the generic trap 6 and specific trap 3 messages, which are sent in parallel with the call home to IBM.

To configure SNMP for the DS8900F, first get the destination address for the SNMP trap and information about the port on which the trap daemon listens.

Standard port: The standard port for SNMP traps is port 162.

12.1.1 Message Information Base file

The DS8900F storage system provides a Message Information Base (MIB) file that describes the SNMP trap objects. Load the file by using the software that is used for enterprise and SNMP monitoring.

The file is in the `snmp` subdirectory on the Data Storage Command-line Interface (DS CLI) installation CD or the DS CLI installation CD image. The image is available at [IBM Fix Central](#).

12.1.2 Predefined SNMP trap requests

An SNMP agent can send SNMP trap requests to SNMP managers to inform them about a change of values or status on the IP host where the agent is running. Seven predefined types of SNMP trap requests exist, as shown in Table 12-1.

Table 12-1 *SNMP trap request types*

Trap type	Value	Description
<code>coldStart</code>	0	Restart after a crash.
<code>warmStart</code>	1	Planned restart.
<code>linkDown</code>	2	Communication link is down.
<code>linkUp</code>	3	Communication link is up.
<code>authenticationFailure</code>	4	Invalid SNMP community string was used.
<code>egpNeighborLoss</code>	5	Exterior Gateway Protocol (EGP) neighbor is down.
<code>enterpriseSpecific</code>	6	Vendor-specific event happened.

Each trap message contains an object identifier (OID) and a value, as shown in Table 12-1 on page 424, to notify you about the cause of the trap message. You can also use type 6, the *enterpriseSpecific* trap type, when you must send messages that do not fit the predefined trap types. For example, the DS8900F uses this type for notifications that are described in this chapter.

12.2 SNMP notifications

The Management Console (MC), which is also known as the Hardware Management Console (HMC), of the DS8900F sends an SNMPv1 trap in the following cases:

- ▶ A serviceable event is reported to IBM by using Call Home.
- ▶ An event occurs in the Copy Services (CS) configuration or processing.
- ▶ When the Global Mirror (GM) operation pauses on the consistency group boundary.
- ▶ When the GM operation fails to unsuspend one or more Global Copy (GC) members.
- ▶ A space efficient repository or an overprovisioned volume reaches a user-defined warning watermark.
- ▶ When a rank reaches I/O saturation.
- ▶ When Encryption Key Management issues an alert that communication between the control unit and one or more Encryption Key Manager (EKM) servers is lost, or reconnected.

A serviceable event is posted as a generic trap 6, specific trap 3 message. The specific trap 3 is the only event that is sent for serviceable events and hardware service-related actions (data offload and remote secure connection). For reporting CS events, generic trap 6 and specific traps 100, 101, 102, 200, 202, 210, 211, 212, 213, 214, 215, 216, 217, 218, 219, 220, 225, or 226 are sent.

Note: Consistency group traps (200 and 201) must be prioritized above all other traps. They must be surfaced in less than 2 seconds from the real-time incident.

12.2.1 Serviceable event that uses specific trap 3

Example 12-1 shows the contents of generic trap 6, specific trap 3. The trap holds the following information:

- ▶ Serial number of the DS8900F
- ▶ Event number that is associated with the manageable events from the HMC
- ▶ Reporting storage facility image (SFI)
- ▶ System reference code (SRC)
- ▶ Location code of the part that is logging the event

The SNMP trap is sent in parallel with a call home for service to IBM and email notification (if configured).

Example 12-1 SNMP special trap 3 of a DS8000

```
Manufacturer=IBM
ReportingMTMS=2107-994*75LAH80
ProbNm=3084
LparName=SF1300960ESS01
FailingEnclosureMTMS=2107-994*75LAH80
```

```
SRC=BE80CB13
EventText=Recovery error,the device hardware error.
FruLoc=Part Number 98Y4317 FRU CCIN U401
FruLoc=Serial Number 1731000A39FC
FruLoc=Location Code U2107.D04.J05S032-P1-D2
```

For open events in the event log, a trap is sent every 8 hours until the event is closed.

12.2.2 Copy Services event traps

For state changes in a remote CS environment, 13 traps are implemented. The 1xx traps are sent for a state change of a physical link connection. The 2xx traps are sent for state changes in the logical CS setup. For all of these events, no call home is generated and IBM is not notified.

This chapter describes only the messages and the circumstances when traps are sent by the DS8900F. For more information about these functions and terms, see *IBM DS8000 Copy Services: Updated for IBM DS8000 Release 9.1*, SG24-8367.

Physical connection events

Within the trap 1xx range, a state change of the physical links is reported. The trap is sent if the physical remote copy link is interrupted. The Link trap is sent from the primary system. The PLink and SLink columns were used only by the 2105 Enterprise Storage Server disk unit.

If one or several links (but not all links) are interrupted, a trap 100 (Example 12-2) is posted. Trap 100 indicates that the redundancy is degraded. The reference code (RC) column in the trap represents the return code for the interruption of the link.

Example 12-2 Trap 100: Remote Mirror and Remote Copy links degraded

```
PPRC Links Degraded
UNIT: Mnf Type-Mod SerialNm LS
PRI:  IBM 2107-981 75-ZA571 12
SEC:  IBM 2107-981 75-CYK71 24
Path: Type PP  PLink SP  SLink RC
1:    FIBRE 0143 XXXXXX 0010 XXXXXX 15
2:    FIBRE 0213 XXXXXX 0140 XXXXXX 0K
```

If all of the links are interrupted, a trap 101 (Example 12-3) is posted. This event indicates that no communication between the primary and the secondary system is possible.

Example 12-3 Trap 101: Remote Mirror and Remote Copy links are inoperable

```
PPRC Links Down
UNIT: Mnf Type-Mod SerialNm LS
PRI:  IBM 2107-981 75-ZA571 10
SEC:  IBM 2107-981 75-CYK71 20
Path: Type PP  PLink SP  SLink RC
1:    FIBRE 0143 XXXXXX 0010 XXXXXX 17
2:    FIBRE 0213 XXXXXX 0140 XXXXXX 17
```

After the DS8900F can communicate again by using any of the links, trap 102 (Example 12-4) is sent after one or more of the interrupted links are available again.

Example 12-4 Trap 102: Remote Mirror and Remote Copy links are operational

```
PPRC Links Up
UNIT: Mnf Type-Mod SerialNm LS
PRI: IBM 2107-981 75-ZA571 21
SEC: IBM 2107-981 75-CYK71 11
Path: Type PP PLink SP SLink RC
1: FIBRE 0010 XXXXXX 0143 XXXXXX OK
2: FIBRE 0140 XXXXXX 0213 XXXXXX OK
```

Remote Mirror and Remote Copy events

If you configured consistency groups, and a volume within this consistency group is suspended because of a write error to the secondary device, trap 200 is sent, as shown in Example 12-5. One trap for each logical subsystem (LSS) that is configured with the consistency group option is sent. This trap can be handled by automation software, such as IBM Copy Services Manager to freeze this consistency group. The SR column in the trap represents the suspension reason (SR) code, which explains the cause of the error that suspended the Remote Mirror (Peer-to-Peer Remote Copy (PPRC)) and Remote Copy group. The SR codes are listed in Table 12-2 on page 430.

Example 12-5 Trap 200: LSS pair consistency group Remote Mirror and Remote Copy pair error

```
LSS-Pair Consistency Group PPRC-Pair Error
UNIT: Mnf Type-Mod SerialNm LS LD SR
PRI: IBM 2107-981 75-ZA571 84 08
SEC: IBM 2107-981 75-CYM31 54 84
```

Trap 202, as shown in Example 12-6, is sent if a Remote Copy pair goes into a suspend state. The trap contains the serial number (SerialNm) of the primary and secondary machine, the LSS (LS), and the logical device (LD). To avoid SNMP trap flooding, the number of SNMP traps for the LSS is throttled. The complete suspended pair information is represented in the summary.

The last row of the trap represents the suspend state for all pairs in the reporting LSS. The suspended pair information contains a hexadecimal string of 64 characters. By converting this hex string into binary code, each bit represents a single device. If the bit is 1, the device is suspended. Otherwise, the device is still in full duplex mode.

Example 12-6 Trap 202: Primary Remote Mirror and Remote Copy devices on LSS suspended due to an error

```
Primary PPRC Devices on LSS Suspended Due to Error
UNIT: Mnf Type-Mod SerialNm LS LD SR
PRI: IBM 2107-981 75-ZA571 28 00 01
SEC: IBM 2107-981 75-CZM21 a8 00
Start: 2015/11/14 10:30:32 CST
PRI Dev Flags (1 bit/Dev, 1=Suspended):
C00000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
```

Trap 210, as shown in Example 12-7, is sent when a consistency group in a GM environment was successfully formed.

Example 12-7 Trap 210: Global Mirror initial consistency group successfully formed

Asynchronous PPRC Initial Consistency Group Successfully Formed
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-ZA571
Session ID: 4002

Trap 211, as shown in Example 12-8, is sent if the GM setup is in a severe error state in which no attempts are made to form a consistency group.

Example 12-8 Trap 211: Global Mirror session is in an unrecoverable state

Asynchronous PPRC Session is in an unrecoverable State
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-CYM21
Session ID: 4002

Trap 212, as shown in Example 12-9, is sent when a consistency group cannot be created in a GM relationship for one of the following reasons:

- ▶ Volumes were taken out of a copy session.
- ▶ The Remote Copy link bandwidth might not be sufficient.
- ▶ The Fibre Channel (FC) link between the primary and secondary system is not available.

Example 12-9 Trap 212: Global Mirror consistency group failure - Retry is attempted

Asynchronous PPRC Consistency Group Failure - Retry will be attempted
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-ZA571
Session ID: 4002

Trap 213, as shown in Example 12-10, is sent when a consistency group in a GM environment can be formed after a previous consistency group formation failure.

Example 12-10 Trap 213: Global Mirror consistency group successful recovery

Asynchronous PPRC Consistency Group Successful Recovery
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-ZA571
Session ID: 4002

Trap 214, as shown in Example 12-11, is sent if a GM session is ended by using the DS CLI `rmgmir` command or the corresponding GUI function.

Example 12-11 Trap 214: Global Mirror master terminated

Asynchronous PPRC Master Terminated
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-ZA571
Session ID: 4002

As shown in Example 12-12, trap 215 is sent if, in the GM environment, the master detects a failure to complete the FlashCopy commit. The trap is sent after many commit retries fail.

Example 12-12 Trap 215: Global Mirror FlashCopy at remote site unsuccessful

Asynchronous PPRC FlashCopy at Remote Site Unsuccessful
A UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-CZM21
Session ID: 4002

Trap 216, as shown in Example 12-13, is sent if a GM master cannot end the GC relationship at one of its subordinates. This error might occur if the master is ended by using the **rmgmir** command but the master cannot end the copy relationship on the subordinate.

You might need to run a **rmgmir** command against the subordinate to prevent any interference with other GM sessions.

Example 12-13 Trap 216: Global Mirror subordinate termination unsuccessful

Asynchronous PPRC subordinate Termination Unsuccessful
UNIT: Mnf Type-Mod SerialNm
Master: IBM 2107-981 75-ZA571
Subordinate: IBM 2107-981 75-CYM31
Session ID: 4002

Trap 217, as shown in Example 12-14, is sent if a GM environment is suspended by the DS CLI command **pausegmir** or the corresponding GUI function.

Example 12-14 Trap 217: Global Mirror paused

Asynchronous PPRC Paused
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-CYM31
Session ID: 4002

Trap 218, as shown in Example 12-15, is sent if a GM exceeded the allowed threshold for failed consistency group formation attempts.

Example 12-15 Trap 218: Global Mirror number of consistency group failures exceeds threshold

Global Mirror number of consistency group failures exceed threshold
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-ZA571
Session ID: 4002

Trap 219, as shown in Example 12-16, is sent if a GM successfully formed a consistency group after one or more formation attempts previously failed.

Example 12-16 Trap 219: Global Mirror first successful consistency group after prior failures

Global Mirror first successful consistency group after prior failures
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-ZA571
Session ID: 4002

Trap 220, as shown in Example 12-17, is sent if a GM exceeded the allowed threshold of failed FlashCopy commit attempts.

Example 12-17 Trap 220: Global Mirror number of FlashCopy commit failures exceeds threshold

```
Global Mirror number of FlashCopy commit failures exceed threshold
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-ZA571
Session ID: 4002
```

Trap 225, as shown in Example 12-18, is sent when a GM operation paused on the consistency group boundary.

Example 12-18 Trap 225: Global Mirror paused on the consistency group boundary

```
Global Mirror operation has paused on the consistency group boundary
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-CYM31
Session ID: 4002
```

Trap 226, in Example 12-19, is sent when a GM operation failed to unsuspend one or more GC members.

Example 12-19 Trap 226: Global Mirror unsuspend members failed

```
Global Mirror operation has failed to unsuspend one or more Global Copy members
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-CYM31
Session ID: 4002
```

Table 12-2 shows the CS SR codes.

Table 12-2 Copy Services suspension reason codes

SR code	Description
03	The host system sent a command to the primary volume of a Remote Mirror and Remote Copy volume pair to suspend copy operations. The host system might specify an immediate suspension or a suspension after the copy completes and the volume pair reaches a full duplex state.
04	The host system sent a command to suspend the copy operations on the secondary volume. During the suspension, the primary volume of the volume pair can still accept updates, but updates are not copied to the secondary volume. The out-of-sync tracks that are created between the volume pair are recorded in the change recording feature of the primary volume.
05	Copy operations between the Remote Mirror and Remote Copy volume pair were suspended by a primary storage unit secondary device status command. This system resource code can be returned only by the secondary volume.
06	Copy operations between the Remote Mirror and Remote Copy volume pair were suspended because of internal conditions in the storage unit. This system resource code can be returned by the control unit of the primary volume or the secondary volume.
07	Copy operations between the Remote Mirror and Remote Copy volume pair were suspended when the auxiliary storage unit notified the primary storage unit of a state change transition to the simplex state. The specified volume pair between the storage units is no longer in a copy relationship.

SR code	Description
08	Copy operations were suspended because the secondary volume became suspended because of internal conditions or errors. This system resource code can be returned only by the primary storage unit.
09	The Remote Mirror and Remote Copy volume pair was suspended when the primary or auxiliary storage unit was restarted or when the power was restored. The paths to the auxiliary storage unit might not be unavailable if the primary storage unit was turned off. If the auxiliary storage unit was turned off, the paths between the storage units are restored automatically, if possible. After the paths are restored, run the mkpprc command to resynchronize the specified volume pairs. Depending on the state of the volume pairs, you might need to run the rmpprc command to delete the volume pairs and run an mkpprc command to reestablish the volume pairs.
0A	The Remote Mirror and Remote Copy pair was suspended because the host issued a command to freeze the Remote Mirror and Remote Copy group. This system resource code can be returned only if a primary volume was queried.

12.2.3 Thin-provisioning SNMP

The DS8900F can trigger two specific SNMP trap alerts that relate to the thin-provisioning feature. The trap is sent out when certain extent pool capacity thresholds are reached, which causes a change in the extent status attribute. A trap is sent under the following conditions:

- ▶ The extent status is not zero (available space is already below threshold) when the first extent space efficient (ESE) volume is configured.
- ▶ ESE volumes are configured in the extent pool.

Example 12-20 shows an example of generated event trap 221.

Example 12-20 Trap 221: Space-efficient repository or overprovisioned volume reached a warning

```
Space Efficient Repository or Over-provisioned Volume has reached a warning
watermark
Unit: Mnf Type-Mod SerialNm
IBM 2107-981 75-ZA571
Volume Type: repository
Reason Code: 1
Extent Pool ID: f2
Percentage Full: 100%
```

Example 12-21 shows an example of generated event trap 223.

Example 12-21 Trap 223: Extent pool capacity reached a warning threshold

```
Extent Pool Capacity Threshold Reached
UNIT: Mnf Type-Mod SerialNm
IBM 2107-981 75-ZA571
Extent Pool ID: P1
Limit: 95%
Threshold: 95%Status: 0
```

12.3 SNMP configuration

The SNMP for the DS8900F is designed to send traps as notifications. The DS8900F does not include an installed SNMP agent that can respond to SNMP polling. Also, the SNMP community name for CS-related traps is fixed and set to `public`.

12.3.1 SNMP preparation

During the planning for the installation (see 6.3.4, “Monitoring DS8900F with the Management Console” on page 184), the IP addresses of the management system are provided for the IBM Systems Service Representative (IBM SSR). This information must be applied by the IBM SSR during the installation. Also, the IBM SSR can configure the HMC to send a notification for every serviceable event or for only those events that call home to IBM.

The network management server that is configured on the HMC receives all of the generic trap 6, specific trap 3 messages, which are sent in parallel with any events that call home to IBM.

The SNMP alerts can contain a combination of a generic and a specific alert trap. The Traps list outlines the explanations for each of the possible combinations of generic and specific alert traps. The format of the SNMP traps, the list, and the errors that are reported by SNMP are available in the “Generic and specific alert traps” section of the Troubleshooting section of the DS8900F [IBM Documentation](#).

SNMP alert traps provide information about problems that the storage unit detects. You or the IBM SSR must perform corrective action for the related problems.

12.3.2 SNMP configuration with the HMC

Clients can configure SNMP alerts by logging in to the DS8900F Service Web User Interface (WUI). The Service WUI can be started from the DS8000 Storage MC (https://<HMC_ip_address>/service) remotely through a web browser. Access the Service MC and log in by using the following client credentials:

- ▶ User ID: `customer`
- ▶ Password: `cust0mer` (default password)

Note: To configure the operation-related traps, use the DS CLI, as shown in 12.3.3, “SNMP configuration with the DS CLI” on page 436.

Complete the following steps to configure SNMP at the HMC.

Important: The two HMCs must be configured separately.

1. Log in to the Service Management section on the HMC, as shown in Figure 12-1.

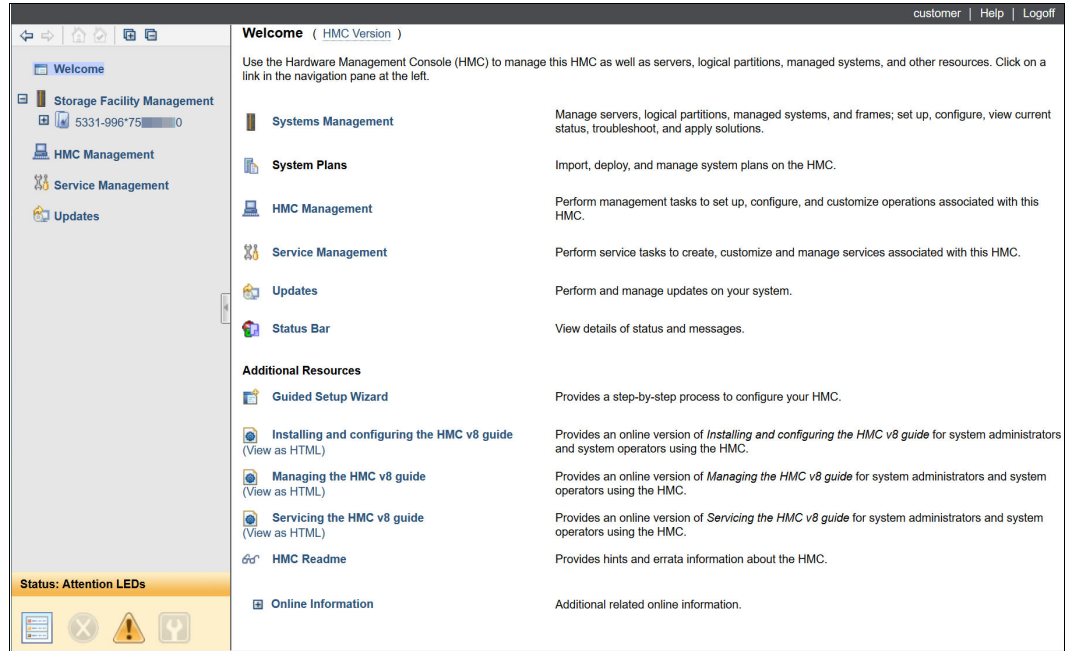


Figure 12-1 HMC Service Management

2. Select **Manage Serviceable Event Notification**, as shown in Figure 12-2.

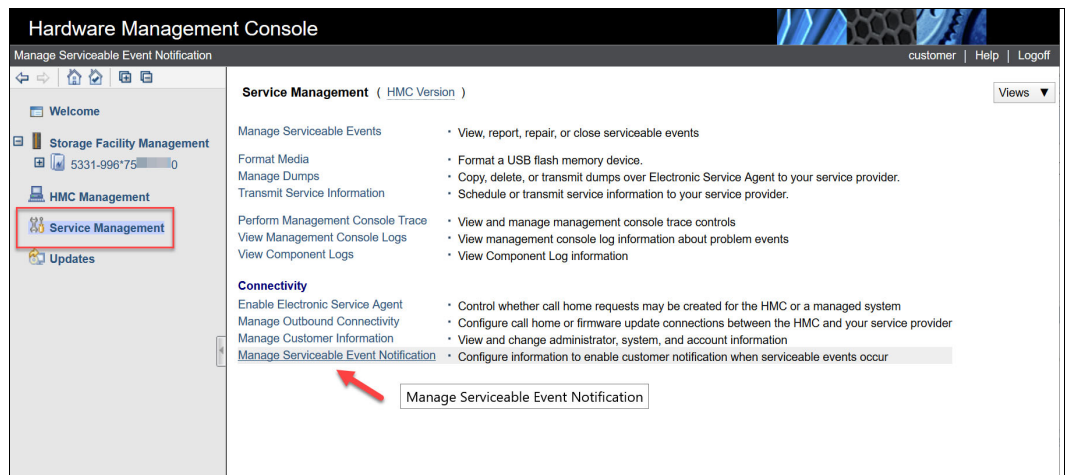


Figure 12-2 HMC Manage Serviceable Event Notification Selection

3. A new window opens. Enter the TCP/IP information of the SNMP server in the SNMP Trap Configuration dialog. Select at least one event out of Trap 3 and Trap 13, as shown in Figure 12-3.

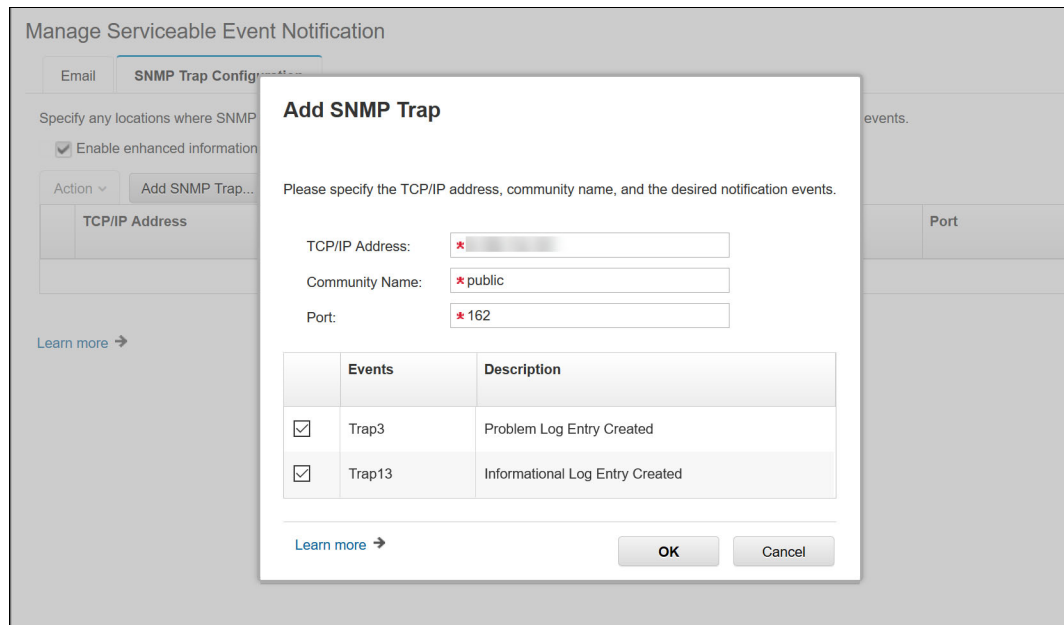


Figure 12-3 HMC Manage Serviceable Event Notification

4. To verify the successful setup of your environment, create a Test Event on your DS8900F MC by selecting the IP address and **Test SNMP Trap**, as shown in Figure 12-4.

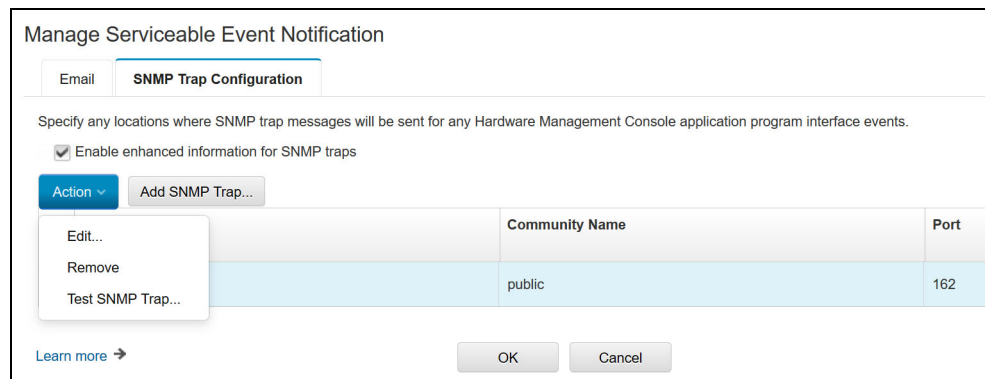


Figure 12-4 HMC test SNMP trap

5. The test result is displayed, as shown in Figure 12-5.




Figure 12-5 HMC test successful

You must check the SNMP server for the successful reception of the test trap.

- The test generates the Service Reference Code BEB20010, and the SNMP server receives the SNMP trap notification, as shown in Figure 12-6.

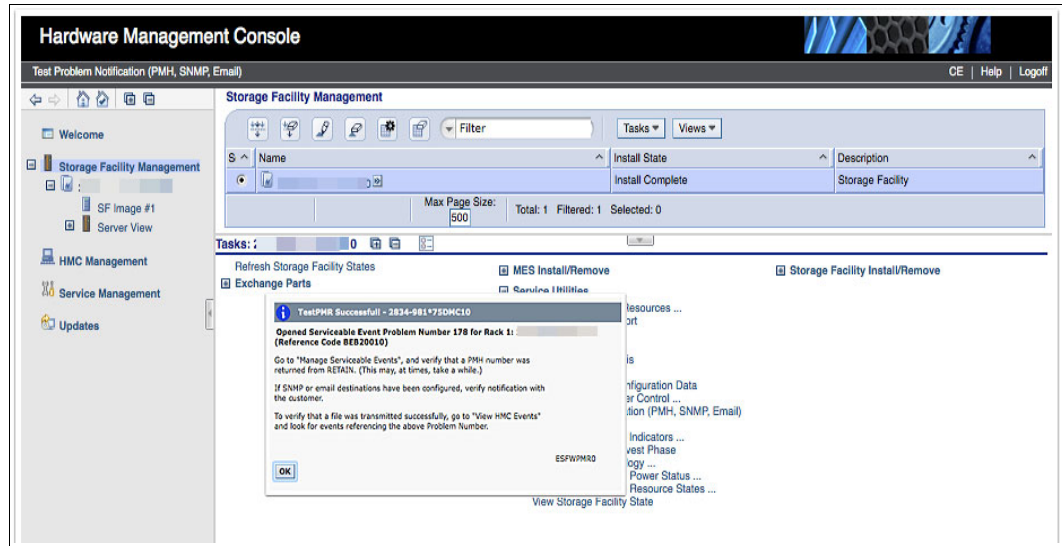


Figure 12-6 HMC SNMP trap test

Errors occurring after configuring the SNMP traps on the HMC are sent to the SNMP server, as shown in Example 12-22.

Example 12-22 Example of an error that is reported to the SNMP system

```

2021-10-11 10:15:22 9.100.123.230(via UDP: [9.100.123.230]:57179->[9.100.121.11])
TRAP, SNMP v1, community public
SNMPv2-SMI::enterprises.2.6.130 Enterprise-Specific Trap (3) Uptime:
14:33:49.15
SNMPv2-SMI::enterprises.2.6.130.3 = STRING: "Oct 11, 2021 10:13:49 AM CEST
Manufacturer=IBM
ReportingMTMS=5331-996*75xxxx0
ProbNm=217
LparName=null
FailingEnclosureMTMS=5331-996*75xxx0
SRC=BC20E504
EventText=Memory subsystem (0x20) reported an error.
Fru1Loc=U78D2.001.WZS08WE-P1-C22-T1-E1
Fru2Loc=U78D2.001.WZS08WE-P1-C22-T1
Fru3Loc=U78D2.001.WZS08WE-P1-C22

PMH=29xxx,075,724
Reporting HMC Hostname=ds8k-r9-xxxx.mainz.de.ibm.com."

```

12.3.3 SNMP configuration with the DS CLI

Perform the configuration process for receiving the operation-related traps, such as for CS, thin provisioning, or encryption, by using the DS CLI. Example 12-23 shows how SNMP is enabled by using the **chsp** command.

Example 12-23 Configuring the SNMP by using the DS CLI

```
dscli> chsp -snmp on -snmpaddr 10.10.10.1,10.10.10.2
CMUC00040I chsp: Storage complex IbmStoragePlex successfully modified.
```

```
dscli> showsp
Name           IbmStoragePlex
desc           -
acct           -
SNMP           Enabled
SNMPaddr       10.10.10.1,10.10.10.2
emailnotify    Disabled
emailaddr      -
emailrelay     Disabled
emailrelayaddr -
emailrelayhost -
numkssupported 4
```

SNMP preparation for the management software

To enable the trap-receiving software to display the correctly decoded message in a human-readable format, load the DS8900F specific Message Information Base file.

The Message Information Base file that is delivered with the latest DS8900F DS CLI CD is compatible with all previous levels of DS8900F Licensed Internal Code (LIC) and previous generations of the DS8000 Product Family. Therefore, ensure that you loaded the latest Message Information Base file that is available.

12.4 Introducing remote support

IBM provides remote support capabilities for the DS8000. Remote support enables storage to communicate with IBM, and allows IBM Support to remotely connect to the system when authorized by the client.

The benefits of remote support are that IBM Support can respond quickly to events that are reported by you or the system.

The following features can be enabled in the DS8900F for remote support:

- ▶ Call Home support (outbound remote support):
 - Reporting problems to IBM
 - Sending heartbeat information
 - Offloading data
- ▶ Remote service (inbound) remote support

IBM Support accesses the DS8900F HMC through a network-based connection.

During the installation and planning phase, complete the remote support worksheets and supply them to the IBM SSR at the time of the installation.

The worksheets have information about your remote support preferences and the network communication requirements that must be fulfilled by your local network.

12.5 IBM policies for remote support

The following guidelines are at the core of the IBM remote support strategies for the DS8000:

- ▶ When the DS8000 transmits service data to IBM, only logs and process memory dumps are gathered for troubleshooting.
- ▶ When a remote session with the DS8000 is needed, the HMC or MC always initiates an outbound connection to predefined IBM servers or ports.
- ▶ IBM maintains multiple-level internal authorizations for any privileged access to the DS8000 components. Only approved IBM Support personnel can gain access to the tools that provide the security codes for HMC CLI access.

Although the MC is based on a Linux operating system (OS), IBM disabled or removed all unnecessary services, processes, and IDs, including standard internet services, such as Telnet (the Telnet server is disabled on the HMC), File Transfer Protocol (FTP), `r` commands (Berkeley `r`-commands and Remote Procedure Call (RPC) commands), and RPC programs.

12.6 Remote support advantages

The following benefits can be realized when you enable remote support on the DS8900F:

- ▶ Serviceable events with related problem data are reported to IBM automatically, and a support call is opened.
- ▶ IBM support personnel can start data analysis and problem isolation immediately, which can reduce the overall time that is required to fix a problem.
- ▶ If more service data is needed, IBM Support can connect to the MC and offload the data for the next level of support.
- ▶ Remote support helps clients to maintain the highest availability of their data.

12.7 Remote support and Call Home

This section details the Call Home characteristics.

12.7.1 Call Home and heartbeat: Outbound

This section describes the Call Home and heartbeat capabilities.

Call Home

Call Home is the capability of the MC to report serviceable events to IBM. The MC also transmits machine-reported product data (MRPD) information to IBM through Call Home. The MRPD information includes installed hardware, configurations, and features. Call Home is configured by the IBM SSR during the installation of the DS8900F by using the customer worksheets. A test call home is placed after the installation to register the machine and verify the Call Home function.

Heartbeat

The DS8900F also uses the Call Home facility to send proactive *heartbeat* information to IBM. The heartbeat configuration can be set by the IBM SSR to send heartbeat information to the customer (through SNMP and email) and IBM. A *heartbeat* is a small message with basic product information that is sent to IBM to ensure that the Call Home function works.

The heartbeat can be scheduled every 1 - 7 days based on the client's preference. When a scheduled heartbeat fails to transmit, a service call with an action plan to verify that the Call Home function is sent to an IBM SSR. The DS8900F uses an internet connection through Transport Layer Security (TLS), which is also known as Secure Sockets Layer (SSL), for Call Home functions.

12.7.2 Data offload: Outbound

For many DS8900F problem events, such as a hardware component failure, a large amount of diagnostic data is generated. This data can include text and binary log files, firmware information, inventory lists, and timelines. These logs are grouped into collections by the component that generated them or the software service that owns them.

The entire bundle is collected together in a *PEPackage*. A DS8900F *PEPackage* can be large, often exceeding 100 MB. In certain cases, more than one *PEPackage* might be needed to diagnose a problem correctly. In certain cases, the IBM Support Center might need an extra memory dump that is internally created by the DS8900F or manually created through the intervention of an operator.

OnDemand Data Dump: The OnDemand Data Dump (ODD) provides a mechanism that allows the collection of debug data for error scenarios. With ODD, IBM can collect data with no impact to the host I/O after an initial error occurs. ODD can be generated by using the DS CLI command `diagsi -action odd` and then offloaded.

The MC is a focal point for gathering and storing all of the data packages. Therefore, the MC must be accessible if a service action requires the information. The data packages must be offloaded from the MC and sent in to IBM for analysis. The offload is performed through the internet through a TLS connection.

12.7.3 Outbound connection types

This section describes the outbound connection options that are available for Call Home and data offload.

Internet through a TLS connection

The preferred remote support connectivity method is internet TLS for MC to IBM communication. TLS is the encryption protocol that was originally developed as a secured web communication standard. Traffic through a TLS proxy is supported by or without authentication based on the client's proxy server configuration.

When the internet is selected as the outbound connectivity method, the MC uses a TLS connection over the internet to connect to IBM. For more information about IBM TLS remote support, planning, and worksheets, see *IBM DS8900F Introduction and Planning Guide*, SC27-9560.

Using the CLI to export data

The **offloadfile** command provides clients with the ability to export different sets of data files. Data sets include the audit log, the IBM Certified Secure Data Overwrite (SDO) certificate, the IBM Easy Tier summary data, the configuration settings, packages to be used by product support teams, the performance summary, the system summary file, and Easy Tier files. Example 12-24 shows exporting the configuration files.

Note: The **offloadfile** command cannot be run from the embedded DS CLI window.

Example 12-24 DS CLI command to download configuration files

```
dsccli> offloadfile -config c:\dsccli_offload
Date/Time: 10 February 2023 18:42:59 CET IBM DSCLI Version: 7.9.30.154 DS:
IBM.2107-75HAL91
CMUC00428I offloadfile: The config file has been offloaded to
c:\dsccli_offload\productSettings.
CMUC00428I offloadfile: The config file has been offloaded to
c:\dsccli_offload\ICSInstall.history.
dsccli>
```

12.8 Remote Support Access (inbound)

IBM took many necessary steps to provide secure network access for the MC. The client can define how and when the IBM SSR can connect to the MC. When remote support access is configured, IBM Support can connect to the MC to start problem analysis and data gathering. This process enables data to be analyzed as fast as possible with an action plan that is created for an onsite IBM SSR, if needed.

Having inbound access that is enabled can greatly reduce the problem resolution time by not waiting for the IBM SSR to arrive onsite to gather problem data and upload it to IBM. With the DS8900, the following inbound connectivity options are available to the client:

- ▶ External Assist On-site (AOS) Gateway
- ▶ Embedded remote access feature

The remote support access connection cannot be used to send support data to IBM.

The support data offload always uses the Call Home feature.

12.8.1 Assist On-site

AOS is an IBM remote access solution. The DS8000 support uses the port-forwarding feature to maintain the DS8000 with an IP-based maintenance tool.

IBM Support encourages you to use AOS as your remote access method.

The remote access connection is secured with TLS 1.2. In addition, a mechanism is implemented so that the HMC communicates only as an outbound connection, but you must specifically allow IBM to connect to the HMC. You can compare this function to a modem that picks up incoming calls. The DS8900F documentation refers to this situation as an *unattended service*.

The connection is always under the control of the DS8000 administrator. Any DS8000 administrator can start and stop the AOS connection.

For more information, see 12.8.4, “Support access management through the DS CLI and DS GUI” on page 441.

When you prefer to have a centralized access point for IBM Support, then an AOS Gateway might be the correct solution. With the AOS Gateway, you install the AOS software externally to a DS8900F HMC. You must install the AOS software on a system that you provide and maintain. IBM Support provides only the AOS software package. Through port-forwarding on an AOS Gateway, you can configure remote access to one or more DS8900F systems or other IBM storage systems.

A simple AOS connection to the DS8000 is shown in Figure 12-7. For more information about AOS, prerequisites, and installation, see *IBM Assist On-site for Storage Overview*, REDP-4889.

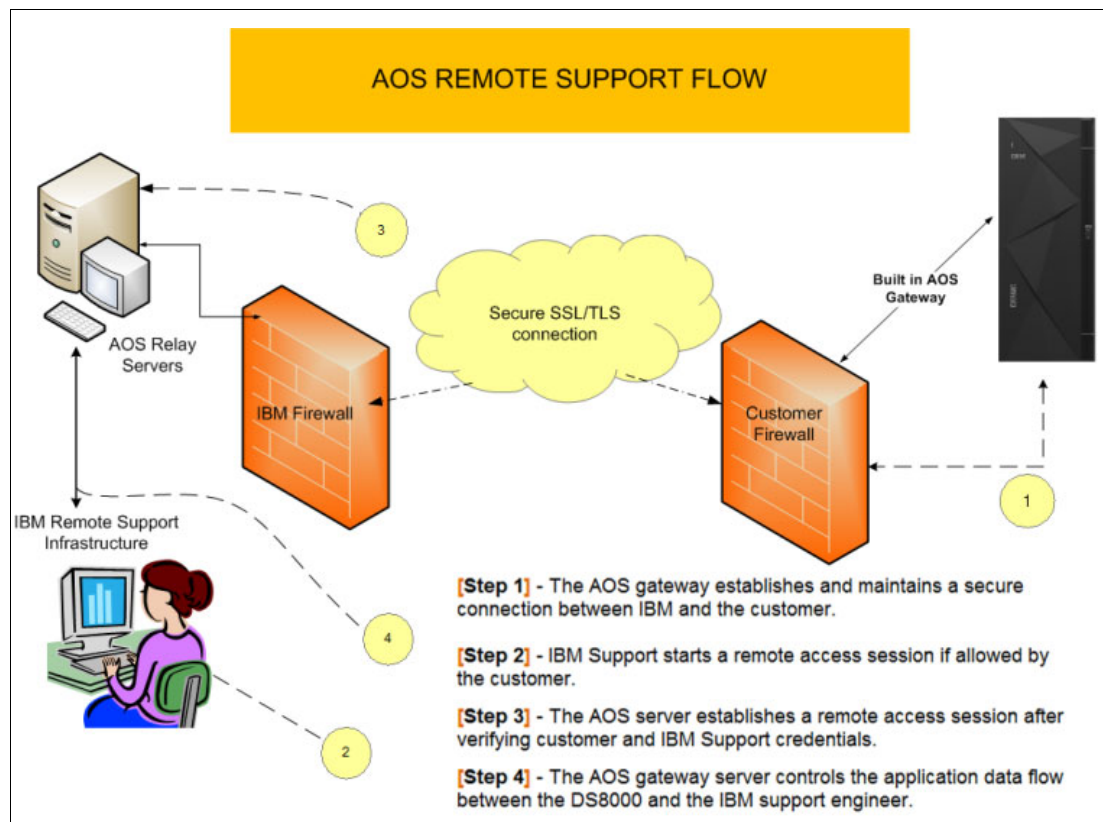


Figure 12-7 DS8000 AOS connection

12.8.2 DS8900F-embedded AOS

AOS is an embedded feature on DS8900F systems. The AOS software package is preinstalled and customized on the MC. This technique eliminates the need to provide an extra system to operate an AOS Gateway. Embedded AOS is a secure and fast broadband form of remote access. You can choose to allow unattended or attended remote access sessions. If you select attended remote access sessions, IBM Support contacts you or the storage operator to start the support session through DS CLI or the DS GUI.

The IBM SSR configures AOS during the installation or a later point by entering information that is provided by the inbound remote support worksheet. The worksheets can be found in *IBM DS8900F Introduction and Planning Guide*, SC27-9560, or in the “Planning” section of the DS8900F [IBM Documentation](#).

In addition, your firewall must allow outbound traffic from the HMC to the AOS infrastructure. The inbound remote support worksheet provides information about the required firewall changes.

For more information about AOS, see *IBM Assist On-site for Storage Overview*, REDP-4889.

12.8.3 IBM Remote Support Center for DS8900F

The HMC is IBM Remote Support Center (RSC) ready. The RSC relies on a single outgoing TCP connection, and it cannot receive inbound connections. Instead of TLS, RSC uses Secure Shell (SSH).

Access to the DS8000 by using RSC is controlled by using either the DS GUI or DS CLI. For more information about RSC, contact your IBM SSR.

12.8.4 Support access management through the DS CLI and DS GUI

All support connections can be enabled or disabled through the DS GUI or DS CLI. The following interfaces can be controlled:

- ▶ The web-based user interface for the IBM SSR on the HMC
- ▶ The SSH CLI access through the local or internal network
- ▶ Remote access through AOS

Using the DS GUI to manage service access

You can control the all service access through the DS Storage Manager GUI through the Access window, which can be opened by clicking **Settings** → **System** → **Advanced**, as shown in Figure 12-8.

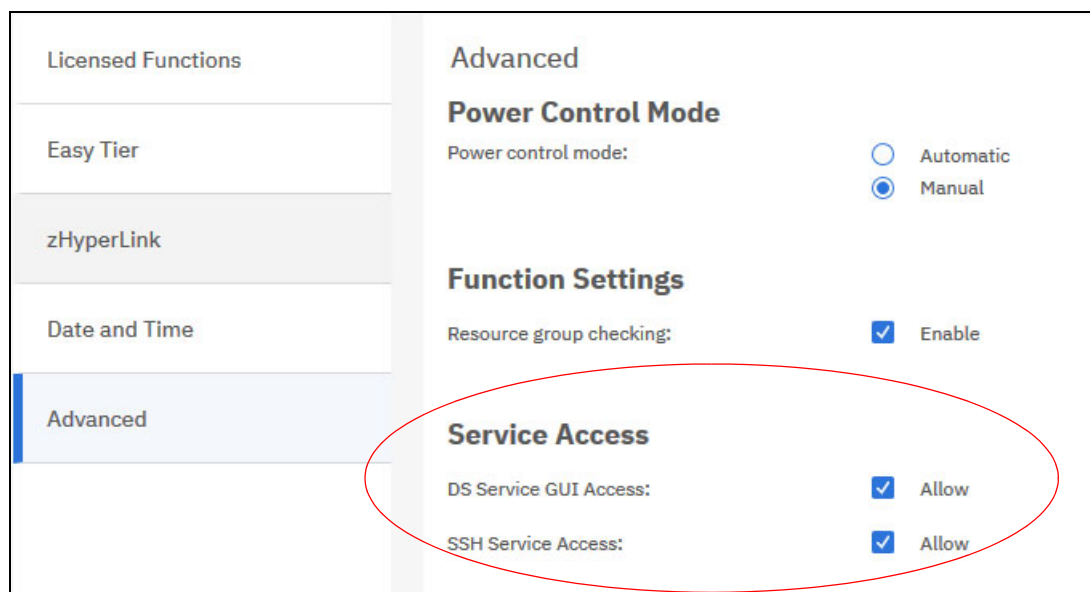


Figure 12-8 Controlling service access through the DS Storage Manager GUI

The following options are available in the Service Access section:

- ▶ **DS Service GUI Access**
Allows authorized IBM SSRs to access the DS Service GUI.
- ▶ **SSH Service Access**
Allows authorized IBM SSRs to access the SSH CLI on the HMC.

Using the DS GUI to manage an embedded AOS connection

You can manage the AOS connection by selecting **Settings** → **Support** → **Assist On-Site**. You can select among three options for the AOS connection: **Start**, **Stop**, or **Restart**. After making the selection, click **Save** in the upper right of the Storage Manager GUI, as shown in Figure 12-9.

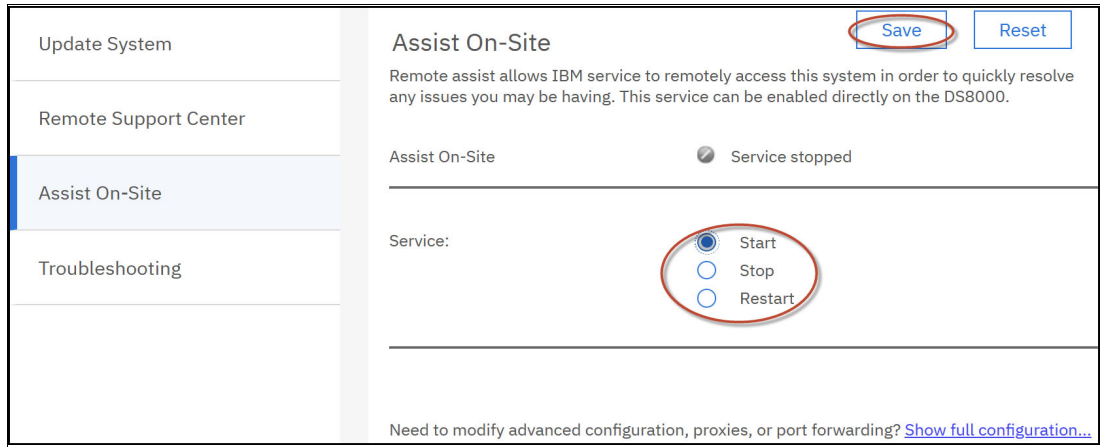


Figure 12-9 Managing the Assist On-Site connection

Using the DS GUI to manage the IBM Remote Support Center connection

To manage the IBM Remote Support Center connection, select **Settings** → **Support** → **Remote Support Center**. You can select among three options: **Open always**, **Open for 2 hours**, and **Closed**. Also, you can generate an access code by selecting the **Access code** checkbox. That code should be provided to the IBM Remote Support Personnel that attempt to connect to the storage system. Figure 12-10 shows the Remote Support Center window.

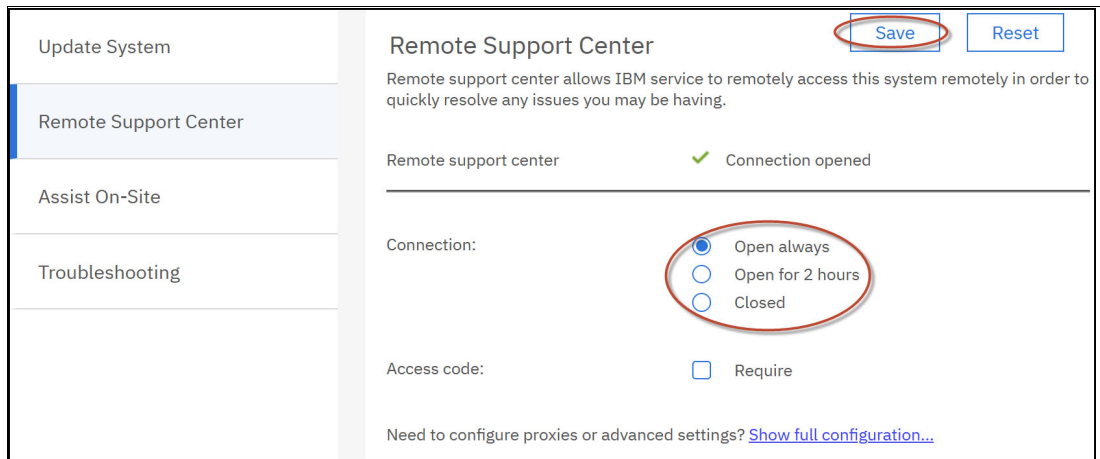


Figure 12-10 Managing the Remote Support Center connection

Using DS CLI to manage service access

You can manage service access to the DS8900F by using DS CLI commands. The following user access security commands are available:

- ▶ **manageaccess:** This command manages the security protocol access settings of an MC for all communications to and from the DS8000 system. You can also use the **manageaccess** command to start or stop outbound virtual private network (VPN) connections instead of using the **setvpn** command.
- ▶ **chaccess:** The command changes one or more access settings of an HMC. Only users with administrator authority can access this command. See the command output in Example 12-25.

```
chaccess [-commandline enable | disable] [-wui enable | disable] [-aos enable | disable] [-rsc enable | disable] -hmc 1|2|all
```

Example 12-25 Output of the chaccess command

```
dsccli> chaccess -aos enable -hmc 1
Date/Time: 10 February 2023 18:34:56 CET IBM DSCLI Version: 7.9.30.154 DS: -
CMUC00441I chaccess: hmc1: The access settings were successfully modified.
dsccli>
```

- ▶ **lsaccess:** This command displays the access settings of the primary and backup MCs:

```
lsaccess [-hmc 1|2|all]
```

See the output in Example 12-26.

Example 12-26 lsaccess command output

```
dsccli> lsaccess -hmc all -l
Date/Time: 10 February 2023 18:36:32 CET IBM DSCLI Version: 7.9.30.154 DS: -
hmc  cmdline wui      modem   cim      aos      rsc      vpn
=====
hmc1 enabled enabled disabled disabled enabled  disabled disabled
hmc2 enabled enabled disabled disabled disabled enabled  disabled
dsccli>
```

Important: The `hmc1` value specifies the primary HMC, and the `hmc2` value specifies the secondary HMC, regardless of how `-hmc 1` and `-hmc 2` were specified during DS CLI start. A DS CLI connection might succeed even if a user inadvertently specifies a primary HMC by using `-hmc 2` and the secondary backup HMC by using `-hmc 1` at DS CLI start.

Client notification of remote login

The MC code records all remote access in a log file. A client can use a DS CLI function to offload this file for audit purposes. The DS CLI function combines the log file that contains all service login information with an IBM Enterprise Storage Server Network Interface (IBM ESSNI) server audit log file that contains all client user login information to provide the client with a complete audit trail of remote access to an MC.

This on-demand audit log mechanism is sufficient for client security requirements for HMC remote access notification.

In addition to the audit log, email notifications and SNMP traps also can be configured at the MC to send notifications in a remote support connection.

12.9 Call Home and Assist On-Site customer-provided certificates

Starting with DS8000 Release 9.3.2, customers can use a customer-provided certificate for both Call Home and for Assist-On-Site (AOS). This enhancement applies only when a customer proxy is used. The certificate must be in the .PEM format, which is the only format that is currently supported.

Before this enhancement, a default certificate was the only option, even if a customer proxy was configured in the communication path between the customer's HMC client and the IBM Electronic Customer Care or IBM AOS servers.

Figure 12-11 shows an overview of the communication path for Call Home with customer proxy configured and using a certificate provided by the customer.

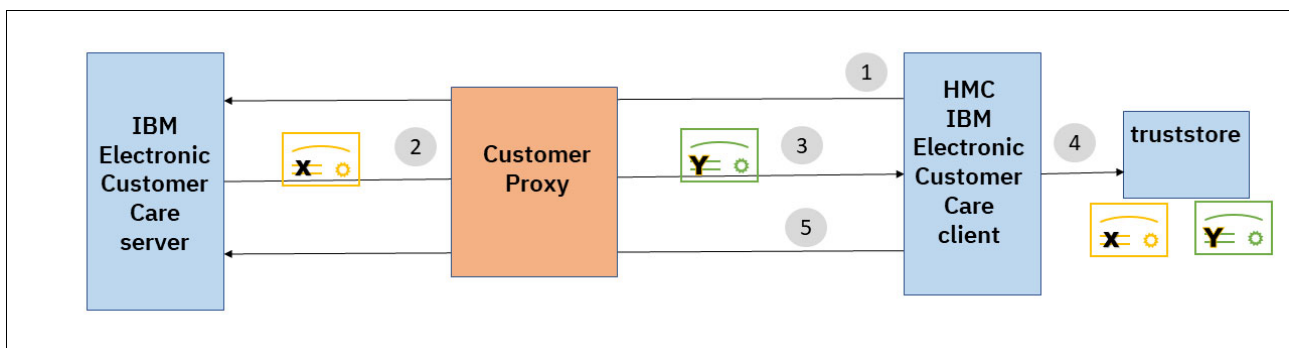


Figure 12-11 Call Home customer-provided certificate - Communication path

1. The HMC IBM Electronic Customer Care client initiates a call home, and the request is sent through the customer proxy to the IBM Electronic Customer Care server.
2. The IBM Electronic Customer Care server always presents the default certificate "X".
3. The customer proxy presents the customer-provided certificate "Y" (instead of the default certificate "X") to the client.
4. The client has a matching certificate ("Y") in the truststore, so the proxy allows communication.

Note: The client still has the original default certificate "X" in the truststore, which is used if Call Home is not configured with customer proxy.

5. The communication continues with the session certificates through the customer proxy.

A similar communication path takes place between the HMC IBM AOS client and the IBM AOS server. The only difference is that the customer-provided certificate "Y" is stored in a property file rather than in a truststore.

Figure 12-12 on page 445 shows an overview of the communication path for AOS with customer proxy configured and using a certificate provided by the customer.

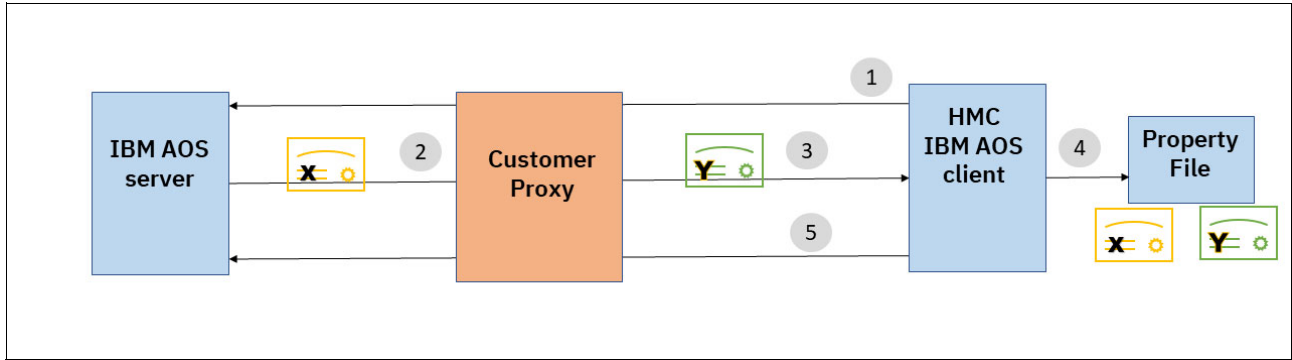


Figure 12-12 AOS customer-provided certificate - Communication path

To install a customer-provided certificate for the Call Home connection, from the DS GUI access **Settings** → **Notifications** → **Call Home** → select **Configure HTTP proxy** → click **Install TLS certificate**. Figure 12-13 shows these settings in the DS GUI. For detailed configuration steps, see [Call home settings](#).

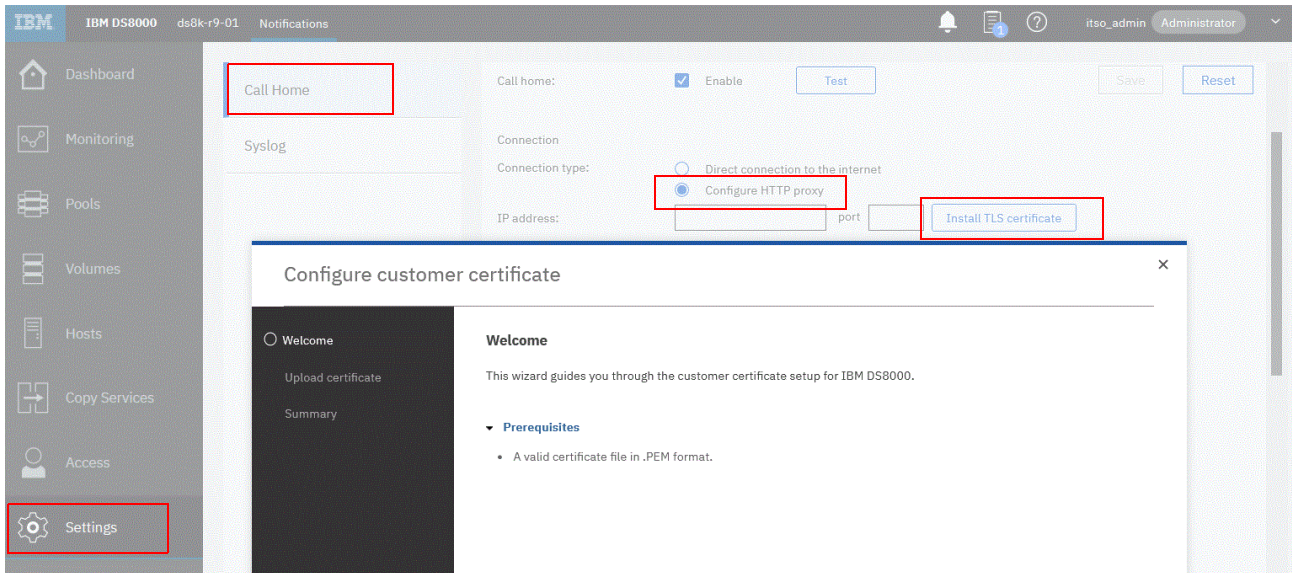


Figure 12-13 Installing a customer-provided certificate for Call Home

To install a customer-provided certificate for the AOS connection, from the DS GUI access **Settings** → **Notifications** → **Assist On-Site** → select **Configure HTTP proxy** → click **Install TLS certificate**. Figure 12-14 on page 446 shows these settings in the DS GUI. For detailed configuration steps, see [Assist On-Site](#).

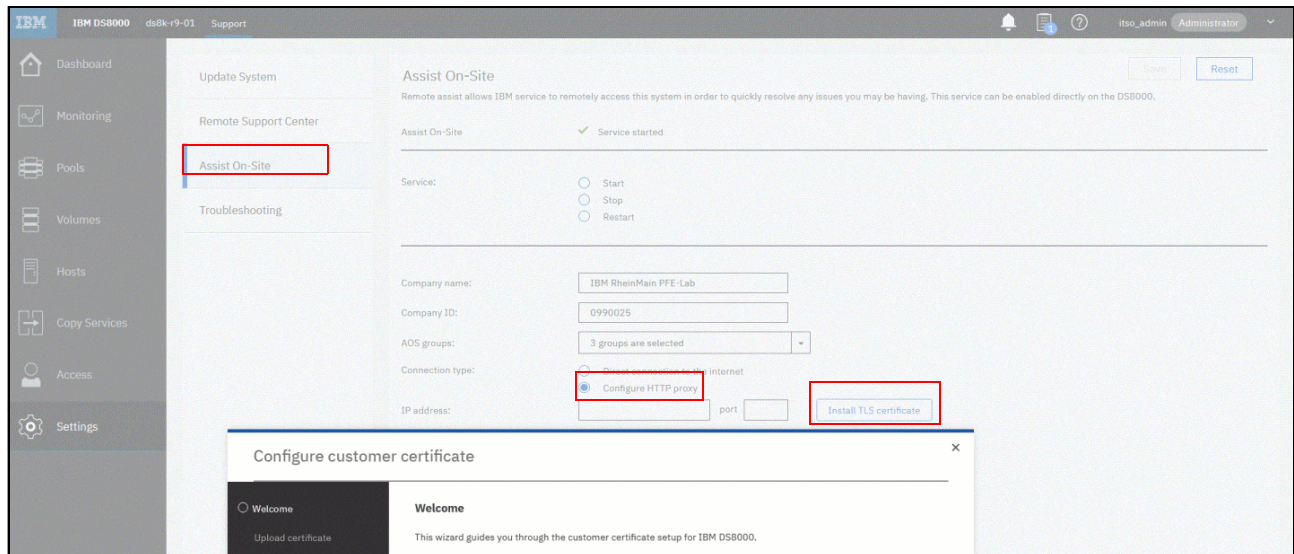


Figure 12-14 Installing a customer-provided certificate for AOS

Guidance from IBM support is available to help setting up this configuration.

12.10 Audit logging

The DS8900F offers an *audit log*, which is an unalterable record of all actions and commands that were initiated by users on the storage system through the DS8000 Storage Management GUI, DS CLI, DS Network Interface (DSNI), or IBM Copy Services Manager. An audit log does not include commands that were received from host systems or actions that were completed automatically by the storage system. The audit logs can be exported and downloaded by the DS CLI or Storage Management GUI.

The DS CLI `offloadauditlog` command provides clients with the ability to offload the audit logs to the client's DS CLI workstation into a directory of their choice, as shown in Example 12-27.

Example 12-27 DS CLI command to download audit logs

```
dsccli> offloadauditlog -logaddr smc1 c:\ha190_audit.txt
Date/Time: May 16, 2022 7:45:24 PM EEST IBM DSCLI Version: 7.9.30.154 DS: -
CMUC00243I offloadauditlog: Audit log was successfully offloaded from smc1 to
c:\ha190_audit.txt.
dsccli>
```

The audit log can be exported by using the DS GUI on the Events window by clicking the **Download** icon and then selecting **Export Audit Log**, as shown in Figure 12-15.

The screenshot shows the IBM DS8000 Events page. The left sidebar contains navigation options: Dashboard, Monitoring, Pools, Volumes, Hosts, Copy Services, and Access. The main area displays a table of events with columns for Type, Time, and Description. A dropdown menu is open over the table, showing 'Export Table' and 'Export Audit Log' options.

Type	Time	Description
Processor Node offline	10/08/2019 10:11:41 PM	Processor Node 0 state offline.
Processor Node offline	10/08/2019 10:11:40 PM	Processor Node 1 state offline.
HMC offline	10/08/2019 10:07:23 PM	HMC ds8k-r9-05-hmc2(2) state offline.
HMC offline	10/08/2019 10:05:27 PM	HMC ds8k-r9-05-hmc2(2) state offline.
HMC offline	10/08/2019 10:05:19 PM	HMC ds8k-r9-05-hmc2(2) state offline.
HMC offline	10/08/2019 10:05:18 PM	HMC ds8k-r9-05-hmc2(2) state offline.
Key server failed	10/08/2019 10:00:11 PM	KMIP key server 5696 stat
HMC offline	10/08/2019 09:54:49 PM	HMC ds8k-r9-05(1) state offline.
HMC offline	10/08/2019 09:52:56 PM	HMC ds8k-r9-05(1) state offline.
HMC offline	10/08/2019 09:52:51 PM	HMC ds8k-r9-05(1) state offline.
HMC offline	10/08/2019 08:09:50 PM	HMC ds8k-r9-05(1) state offline.
HMC offline	10/08/2019 08:09:40 PM	HMC ds8k-r9-05(1) state offline.
HMC offline	10/08/2019 06:21:01 PM	HMC ds8k-r9-05-hmc2(2) state offline.

Figure 12-15 Export Audit Log

The downloaded audit log is a text file that provides information about when a remote access session started and ended, and the remote authority level that was applied. A portion of the downloaded file is shown in Example 12-28.

Example 12-28 Audit log entries that relate to a remote support event

```
MST,,1,IBM.2107-75ZA570,N,8036,Authority_to_root,Challenge Key = 'Fy31@C37';
Authority_upgrade_to_root,,,
U,2021/10/02 12:09:49:000
MST,customer,1,IBM.2107-75ZA570,N,8020,WUI_session_started,,,,
U,2021/10/02 13:35:30:000
MST,customer,1,IBM.2107-75ZA570,N,8022,WUI_session_logoff,WUI_session_ended_logged
off,,,
```

The *challenge key* that is presented to the IBM SSR is a part of a two-factor authentication method that is enforced on the MC. It is a token that is shown to the IBM SSR who connects to the DS8900F. The IBM SSR must use the challenge key in an IBM internal system to generate a *response key* that is given to the HMC. The response key acts as a one-time authorization to the features of the HMC. The challenge and response keys change when a remote connection is made.

The challenge-response process must be repeated if the SSR needs higher privileges to access the MC command-line environment. No direct user login and no root login are available on a DS8900F.

Entries are added to the audit file only after the operation completes. All information about the request and its completion status is known. A single entry is used to log request and response information. It is possible, though unlikely, that an operation does not complete because of an operation timeout. In this case, no entry is made in the log.

The audit log entry includes the following information:

- ▶ A log of users that connect or disconnect to the storage manager.
- ▶ A log of user password and user access violations.
- ▶ Many commands that create, remove, or modify the logical configuration, including the command parameters and user ID.

- ▶ A log of commands that modify SFI and storage facility settings, including the command parameters and user ID.
- ▶ A log of CS commands, including command parameters and users.

Note: IBM Copy Services Manager commands are not supported.

Audit logs are automatically trimmed (first in, first out (FIFO)) by the subsystem so that they do not use more than 50 MB of disk storage.

12.11 Using IBM Storage Insights

You can take advantage of more functions that are offered by IBM Storage Insights for Call Home and for events and notifications logging to improve the user experience with DS8900F.

12.11.1 IBM Storage Insights

IBM Storage Insights provides an unparalleled level of visibility across your storage environment to help you manage complex storage infrastructures and make cost-saving decisions. It combines IBM data management leadership with proprietary analytics from IBM Research. As a cloud-based service, it enables you to deploy quickly and save storage administration time while optimizing your storage. It also helps automate aspects of the support process to enable faster resolution of issues.

With the following editions, you can select the capabilities that serve your needs best:

- ▶ IBM Storage Insights
IBM Storage Insights provides a unified view of a storage environment with a diagnostic events feed, an integrated support experience, and key capacity and performance metrics. IBM Storage Insights is available at no cost to owners of IBM block storage systems who sign up for it.
- ▶ IBM Storage Insights Pro
The capacity-based, subscription version is called IBM Storage Insights Pro and includes all the features of IBM Storage Insights plus all functions that you expect from a storage resource management tool, such as more comprehensive views of the performance, capacity, and health of storage resources. It also helps you reduce storage costs and optimize your data center by providing features like intelligent capacity planning, storage reclamation, storage tiering, and advanced performance metrics. The storage systems that you can monitor are expanded to include IBM file, object, software-defined storage (SDS) systems, and non-IBM block and file storage systems.

Combining features such as Call Home, data collectors, a streamlined ticketing process, and proactive support, the problem resolution gains speed so that the stability, capacity, and performance of the DS8900F can be managed more efficiently.

If a problem occurs, you receive help promptly through the unified support experience by completing the following tasks:

- ▶ Open IBM Support tickets for a resource and automatically add a log package to the ticket.
- ▶ Update tickets with a new log package.
- ▶ View the ticket history of open and closed tickets for a device.

A lightweight data collector is installed in your data center to stream performance, capacity, asset, and configuration metadata to your IBM Cloud instance.

The metadata flows in one direction, that is, from your data center to IBM Cloud over HTTPS. In IBM Cloud, your metadata is protected by physical, organizational, access, and security controls.

12.11.2 Getting started with IBM Storage Insights

To start using IBM Storage Insights, complete the following steps:

1. Sign up for IBM Storage Insights at [IBM Storage Insights registration](#) and then complete the following steps:
 - a. Wait for an IBM representative to contact you to get you started, which should occur within 24 hours.
 - b. When you register, specify an owner for IBM Storage Insights. The owner manages access for other users and acts as the main contact.
 - c. You receive a Welcome email when IBM Storage Insights is ready. The email contains a direct link to your dashboard.

2. Install a data collector in your data center to stream performance, capacity, and configuration metadata about storage systems to IBM Storage Insights. Select **Configuration Data Collectors**, click **Deploy Data Collectors**, and get started, as shown in Figure 12-16:
 - a. Choose your preferred OS to download the data collector (Windows, Linux, or AIX).
 - b. Extract the contents of the data collector file on to the virtual machine (VM) or server where you want it to run. 1 GB of memory and 1 GB of disk space are required.
 - c. Run `installDataCollectorService.sh` (Linux or AIX) or `installDataCollectorService.bat` (Windows).

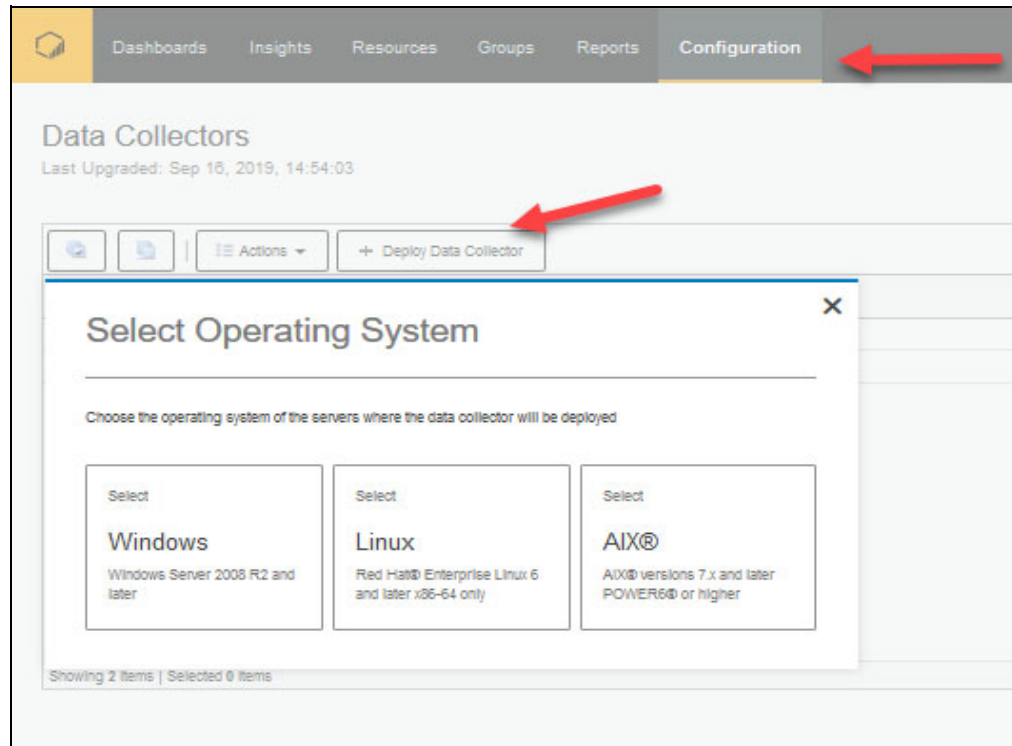


Figure 12-16 Deploy Data Collector

Figure 12-17 shows a system overview where you can access tickets details and the actions that you can take to manage tickets.

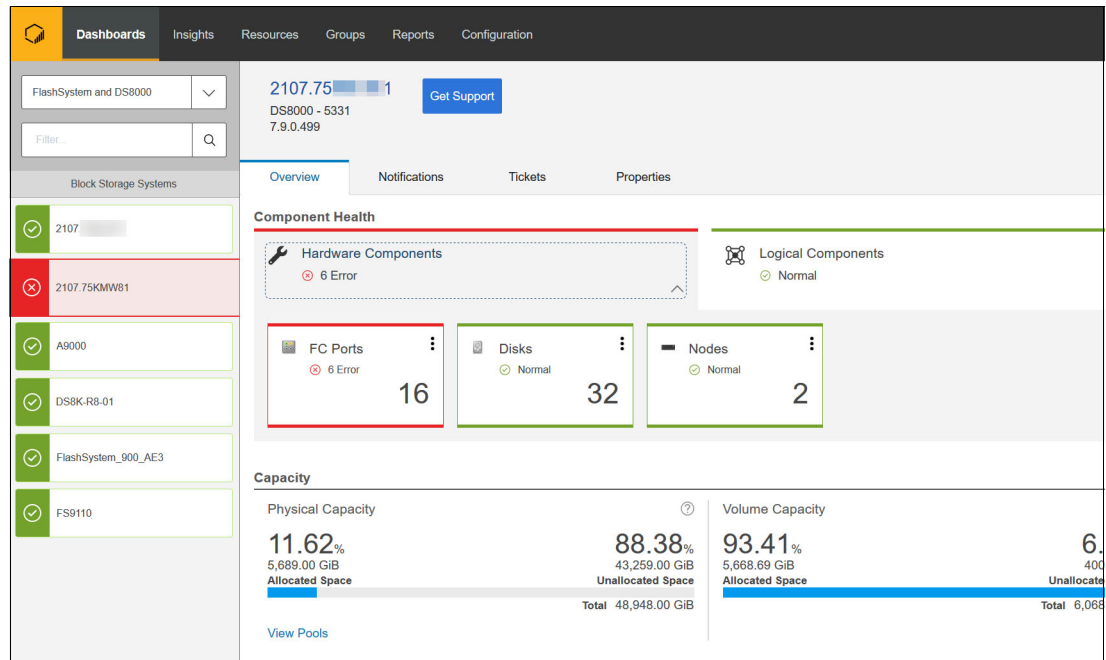


Figure 12-17 System overview

An example of an IBM Storage Insights Pro resources view is the detailed volume information that is shown in Figure 12-18.

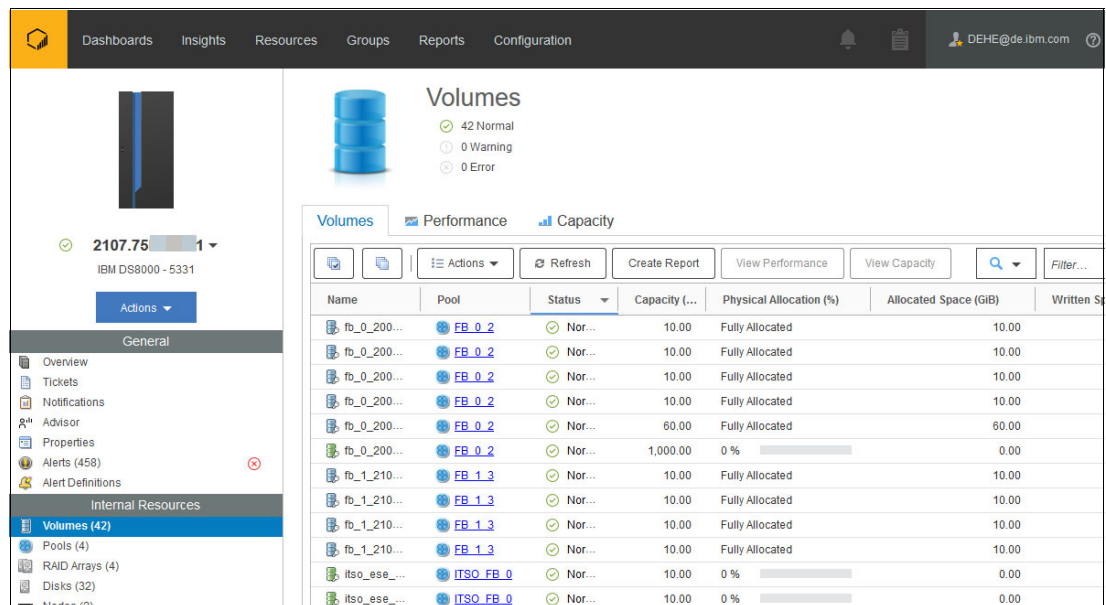


Figure 12-18 IBM Spectrum Insights Pro volume details

Note: If you do not have an IBMid yet, go to [IBM](#) and register.

By using IBM Storage Insights, IBM Remote Support personnel can collect log packages from a device. By default, this feature is not enabled. After the device is added, you must enable the option for IBM Support to collect logs from the device.

To enable this option for IBM Support, select **Configuration** → **Settings** → **IBM Support Log Permission**, and then select **Edit**.

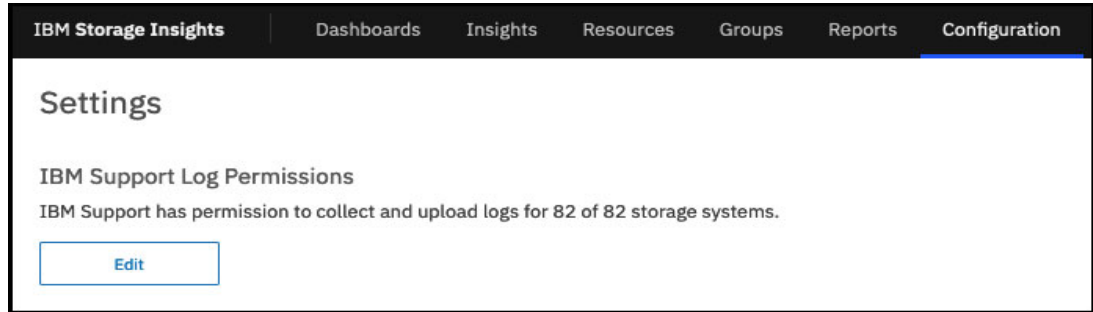


Figure 12-19 IBM Storage Insights: Log Permissions

For more information, see [IBM Storage Insights documentation](#).

12.11.3 Case interaction with IBM Storage Insights

With IBM Storage Insights, a customer can interact and manage each case that any of their DS8000 storage systems opened. By using this tool, you can do the following tasks:

- ▶ Create IBM Support tickets.
- ▶ Update tickets with new log packages or leave a note for IBM Support personnel.
- ▶ View open tickets or the ticket history for each storage system.
- ▶ Interact with the Call Home notifications for each system.

Creating an IBM Support ticket

To create an IBM Support ticket, complete the following steps:

1. Go to the storage system Overview page and select **Actions**, as shown in Figure 12-20. Select **Create/Update Ticket**.

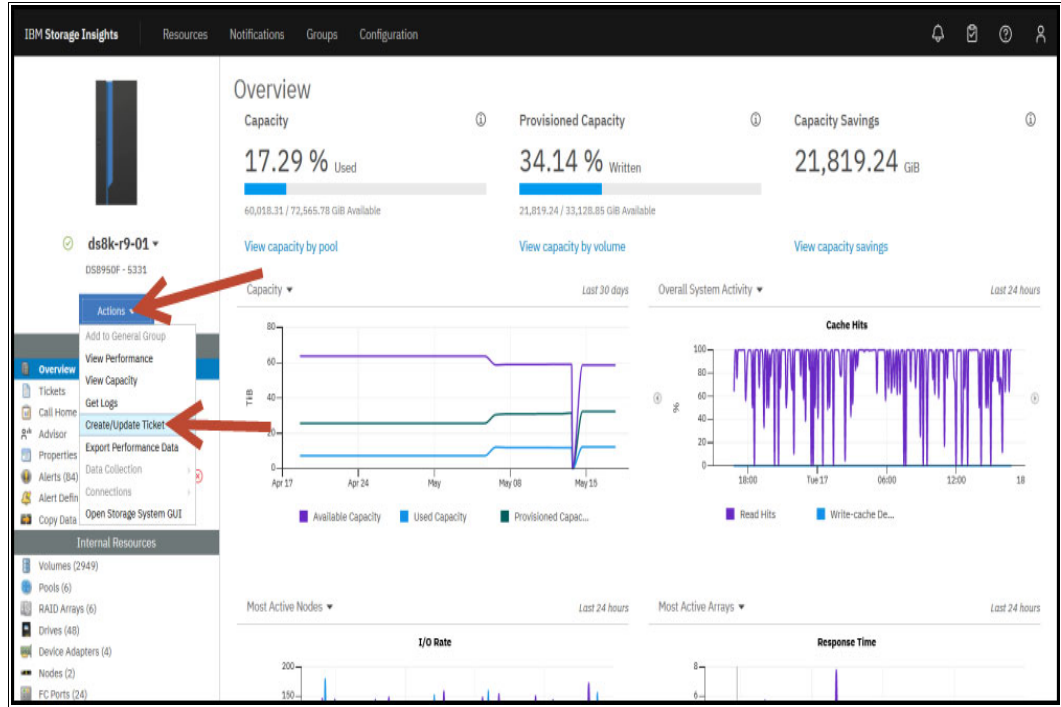


Figure 12-20 IBM Storage Insights: Overview window

2. A new window opens with two options from which to select. Click **Create Ticket**, as shown in Figure 12-21 on page 454.

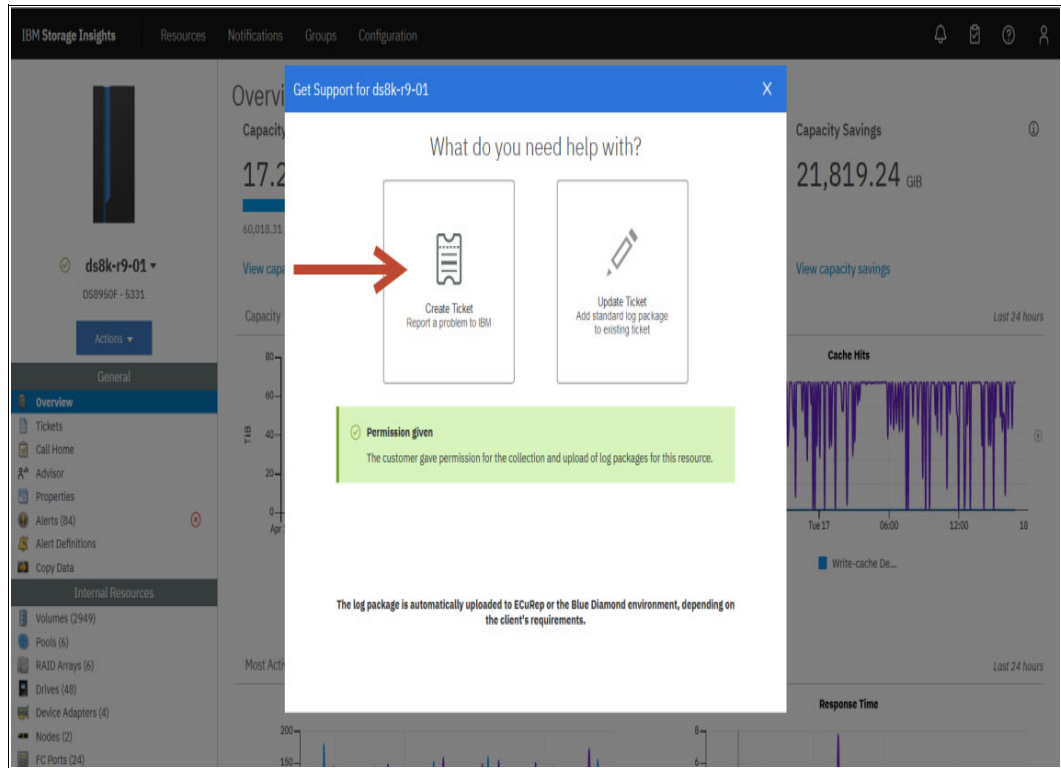


Figure 12-21 IBM Storage Insights: Create Ticket

Figure 12-22 shows the process of collecting the needed ticket information, which includes the details of the DS8000 storage system.

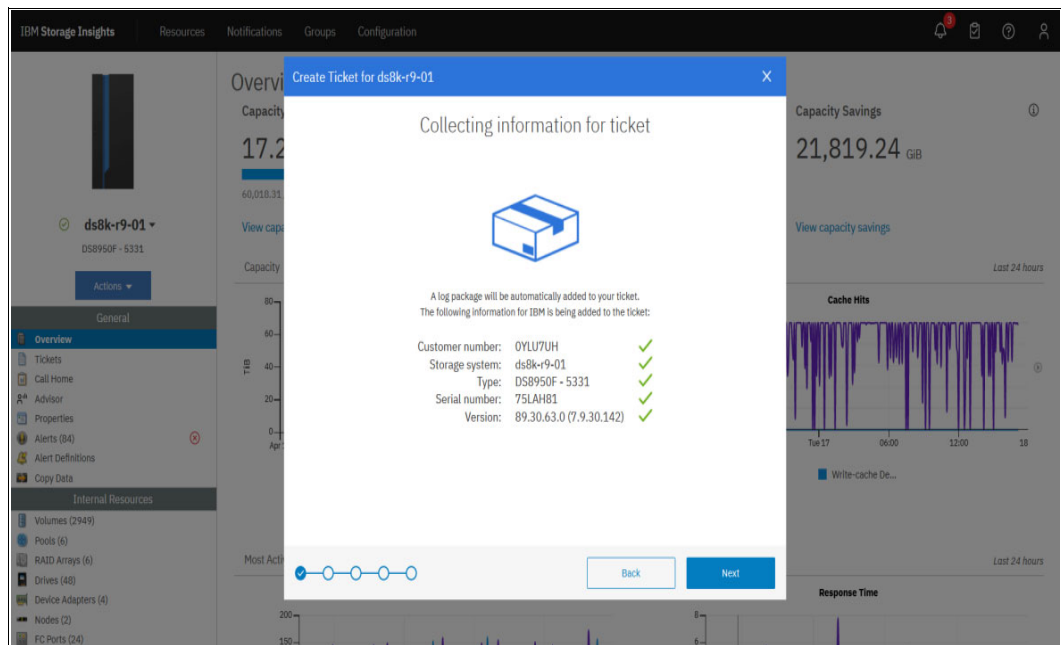


Figure 12-22 IBM Storage Insights: Collecting ticket information

- In Figure 12-22 on page 454, click **Next**. In the window that opens, you describe the reason for opening the ticket with IBM. The mandatory field for the summary is limited to 72 characters. Figure 12-23 shows an optional field that you can use to provide more details.

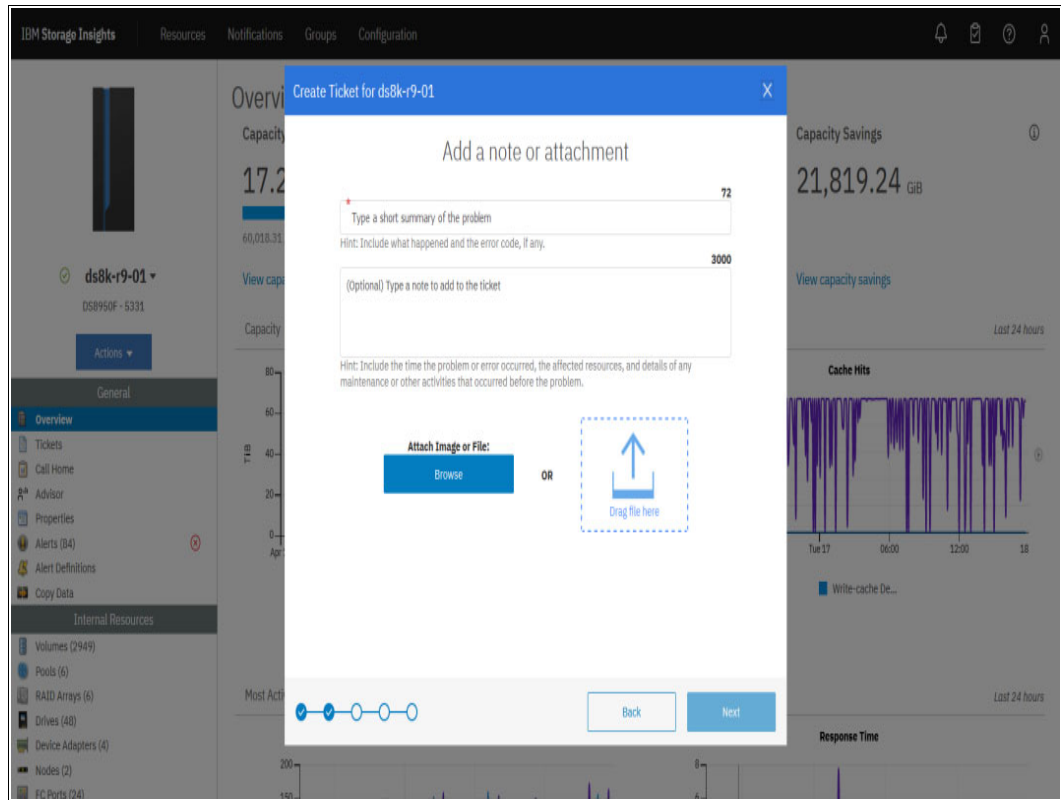


Figure 12-23 IBM Storage Insights: Ticket summary

- Click **Next**. In the window that is shown in Figure 12-24, select the severity of the ticket.

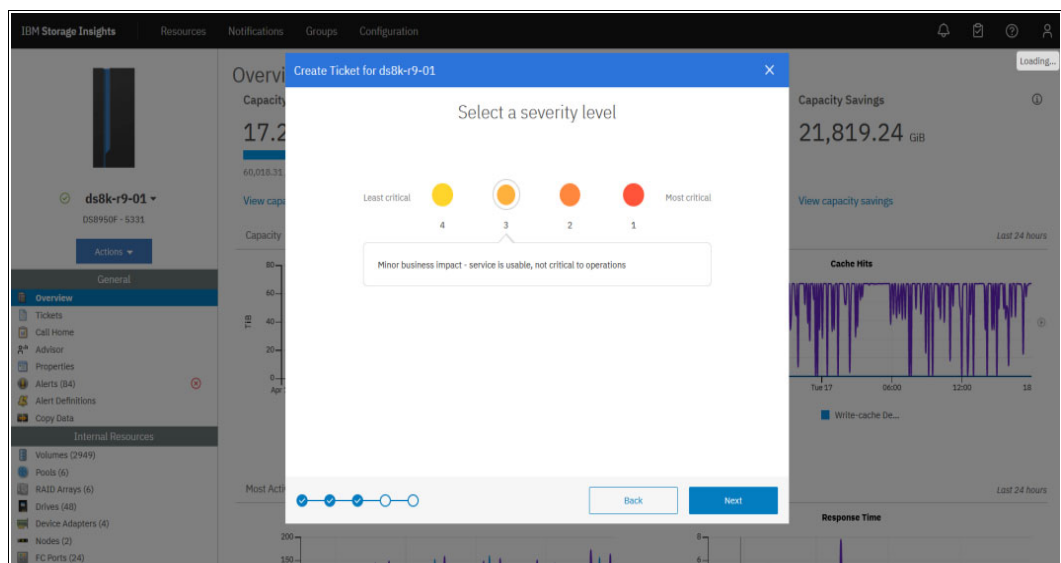


Figure 12-24 IBM Storage Insights: Severity level

5. In the window that is shown in, Figure 12-25, select either **Software/I don't know** or **Hardware**, depending on whether you have a hardware problem or a software problem.

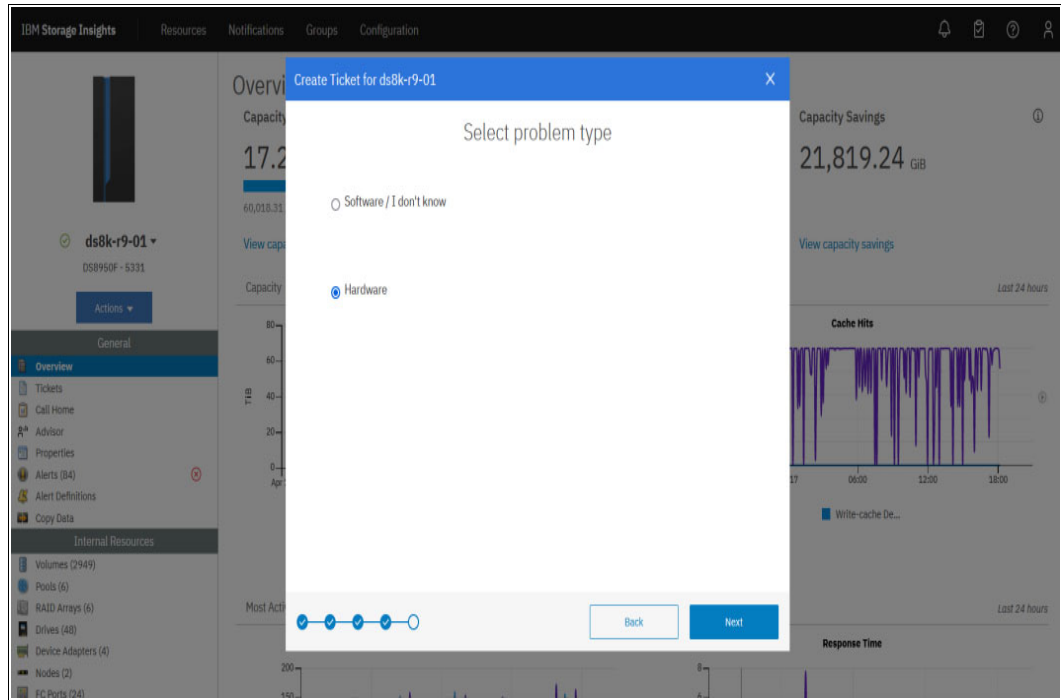


Figure 12-25 IBM Storage Insights: Select problem type

6. Review and verify the details of the IBM ticket that you are about to open, as shown in Figure 12-26. Provide the name of the contact person, along with a valid contact phone number and email address.

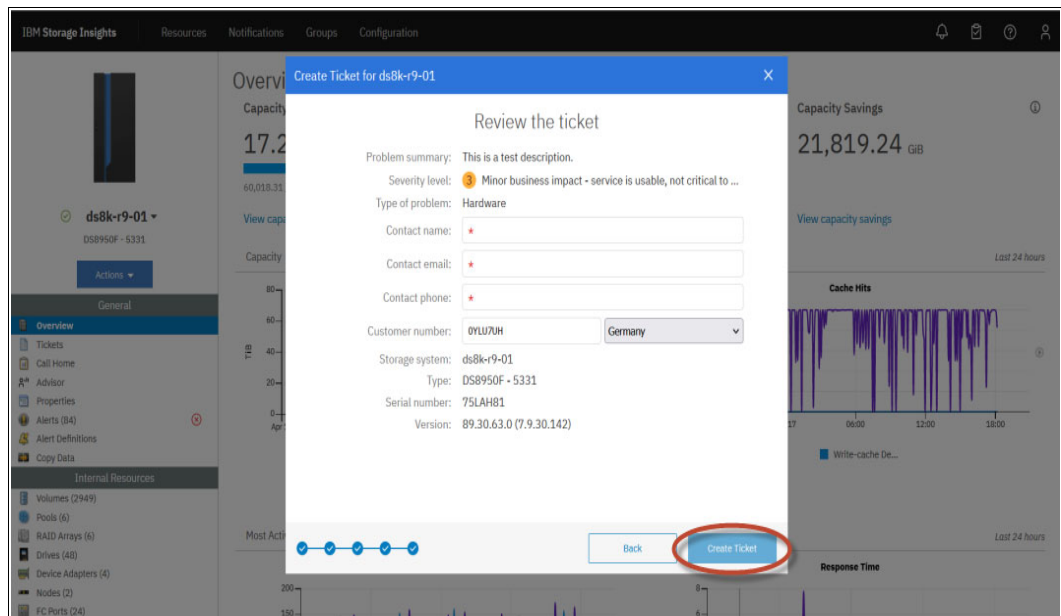


Figure 12-26 IBM Storage Insights: Review the ticket

Updating IBM tickets

Whenever a customer must update a case or wants to add more data to an existing ticket, they can do that task from IBM Storage Insights by completing the following steps:

1. Go to the Overview page for the storage system and select **Actions**, and then select **Create/Update Ticket**, as shown in Figure 12-20 on page 453.
2. In the window that is shown in Figure 12-27, click **Update Ticket**.

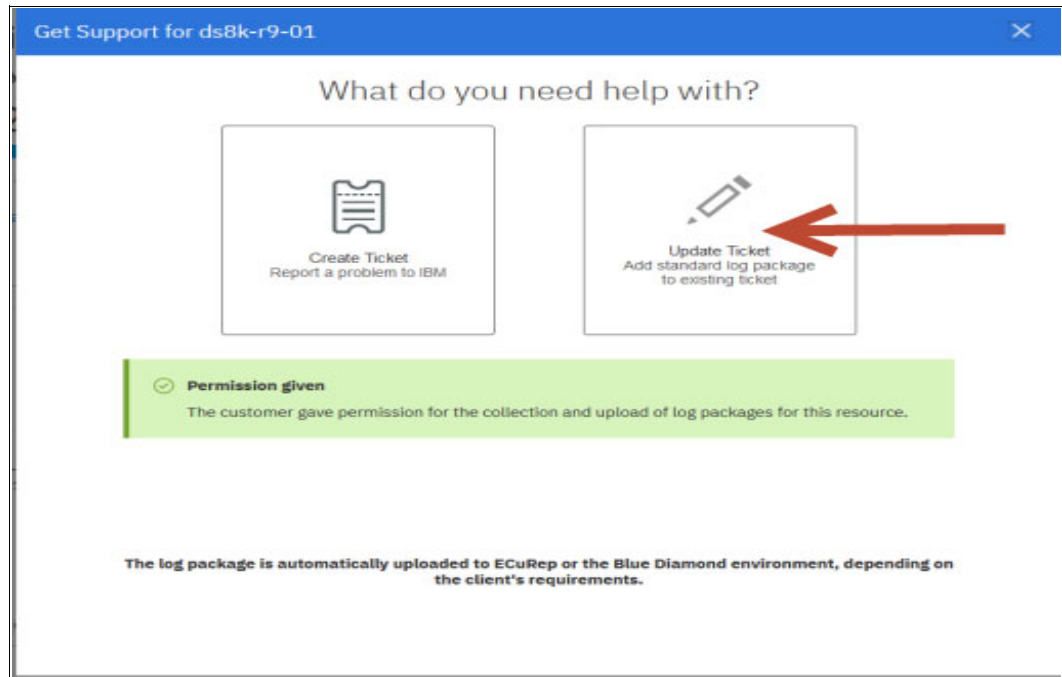


Figure 12-27 IBM Storage Insights: Update Ticket

3. In the window that is shown in Figure 12-28, you can either select one of the open tickets for this machine or type in the ticket number manually.

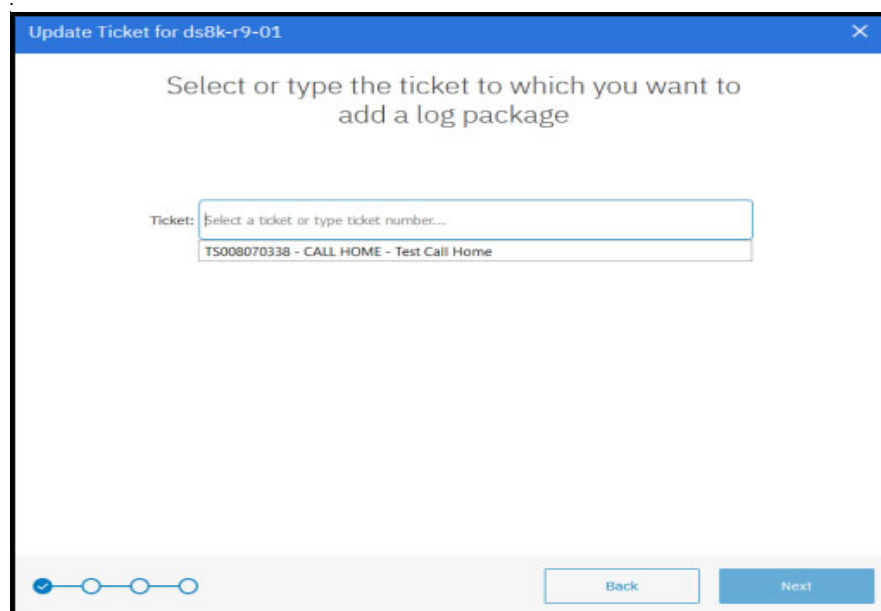


Figure 12-28 IBM Storage Insights: Ticket number

- An automatic process begins after you select the ticket number, which generates a log package.
- In the window that is shown in Figure 12-29, you can leave a message for the IBM Support personnel working on the ticket. Also, you can attach any files that are related to the ticket.

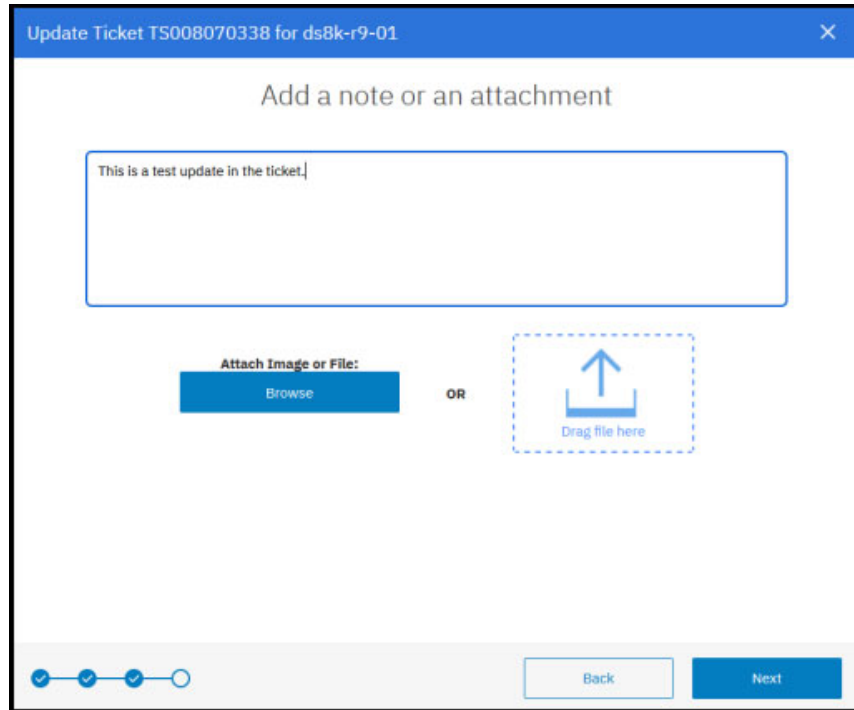


Figure 12-29 IBM Storage Insights: Ticket update and attachment

- In the window that is shown in Figure 12-30, you see a summary of the update that is about to be left in to the ticket. To complete the process, click **Update Ticket**.

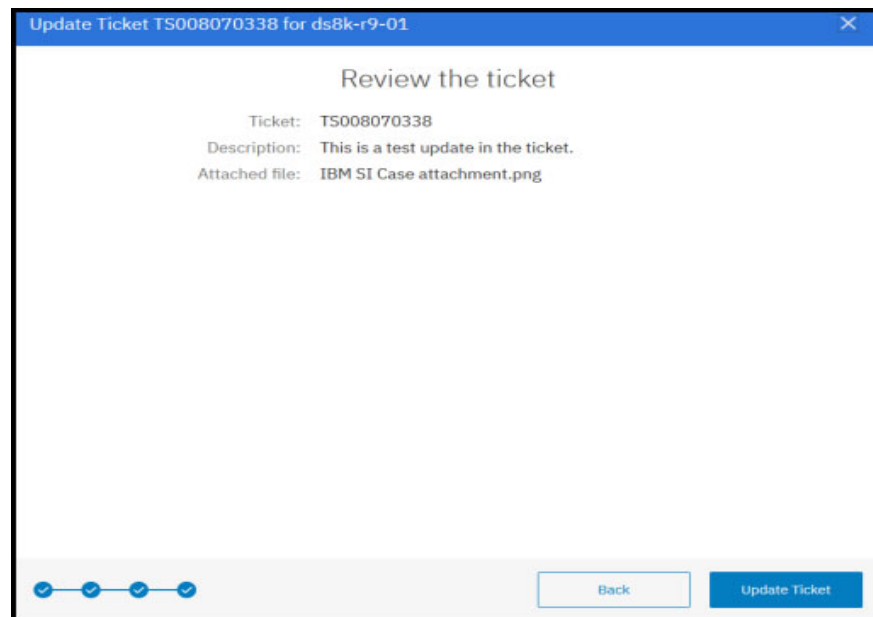


Figure 12-30 IBM Storage Insights: Update summary

12.11.4 IBM Storage Insights: Alert Policies

One of the many features that IBM Storage Insights provides to you is setting Alert Policies for your devices. This function notifies you about certain conditions on a device when that condition is met.

To set alert policies, select **Configuration** → **Alert Policies**, as shown in Figure 12-31.

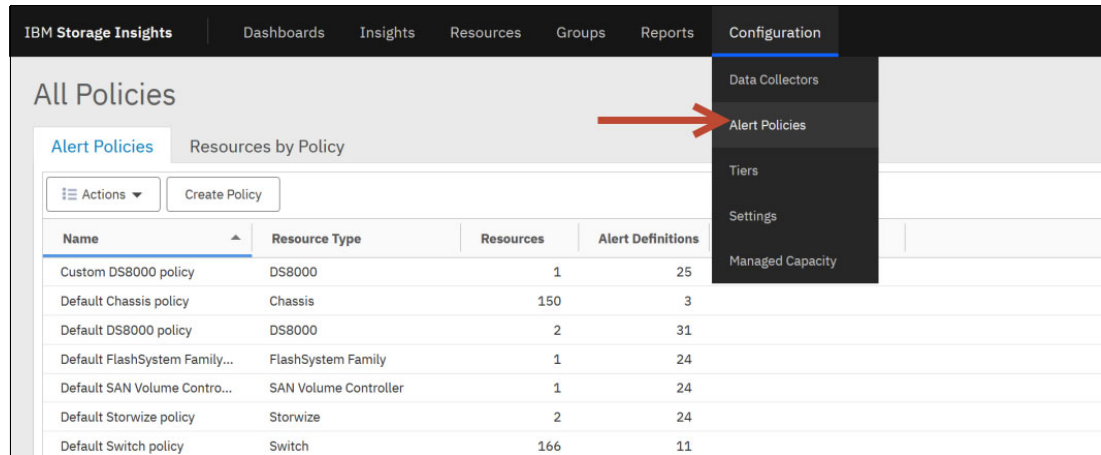


Figure 12-31 IBM Storage Insights: Alert Policies

There are many predefined default policies, but you also can create a policy by selecting **Create Policy**. This action opens a window where you define the policy name, and select the policy type and type of storage system, as shown in Figure 12-32.

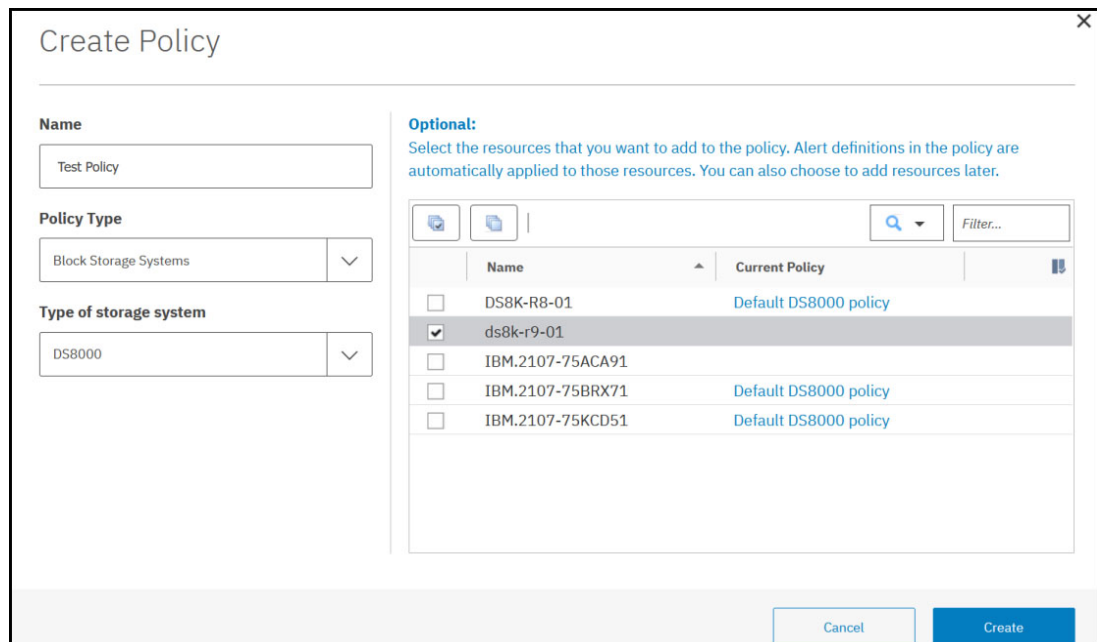


Figure 12-32 IBM Storage Insights: Create Policy

Select **Create**, and the policy is created. Then, define the alert definitions and save the changes, as shown in Figure 12-33 on page 460.

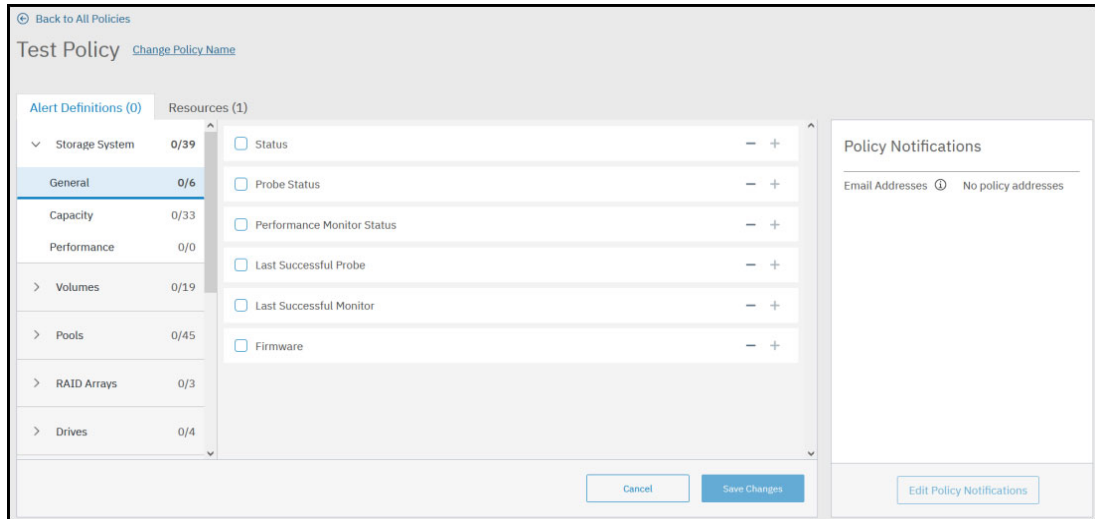


Figure 12-33 IBM Storage Insights: Alert Definitions

12.12 IBM Call Home Connect Cloud

IBM Call Home Connect Cloud is an application that enhances your experience in managing your IBM systems products. It replaces the IBM Call Home Web service.

Some of the benefits of the application are as follows:

- ▶ A single place to manage your assets
- ▶ Access to warranty and maintenance information
- ▶ Real-time updates and notifications
- ▶ Offline mode
- ▶ Access to product-specific tools
- ▶ Software and firmware recommendations

For more information, see

<https://www.ibm.com/support/pages/introducing-call-home-connect-cloud> and <https://www.ibm.com/support/pages/node/6612751>.

IBM Call Home Connect Anywhere

IBM Call Home Connect Anywhere is a mobile application in either Google Play (Android) or the Apple App Store that helps you to monitor your IBM devices from anywhere. Using the application, you can stay informed about any alerts or cases. You can check your device's details, warranty, and contract status.

To work with the application, you must first register for IBM Call Home Connect Cloud. After you can access your IBM Call Home Connect Cloud at

<https://www.ibm.com/support/call-home-connect/cloud/>

register each of your IBM assets by providing its machine type, model, serial number, customer number, and country code. After you complete this task for each device, you can see them in the mobile application.

For more information, see <https://www.ibm.com/support/pages/node/6562485>.

Figure 12-34 shows an example for having various DS8000s and an IBM z15 enlisted with an active DS8910F ticket that is visible on an Apple iOS mobile device when using the IBM Call Home Connect Anywhere application.

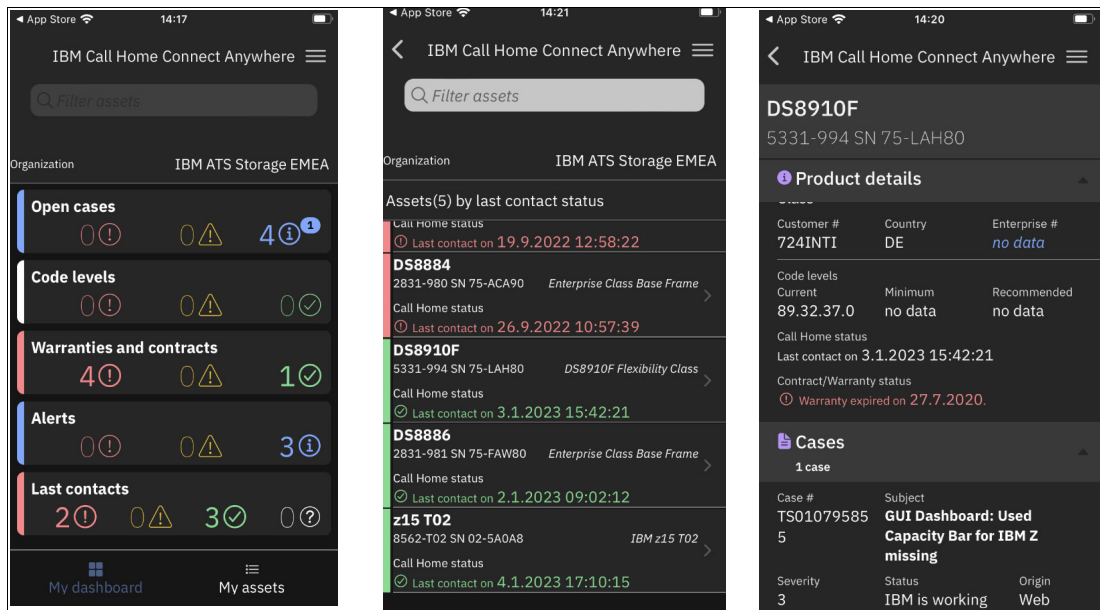


Figure 12-34 IBM Call Home Connect Anywhere: iOS mobile device showing DS8000s and a DS8910F ticket

Abbreviations and acronyms

AC	active	DRP	disaster recovery planning
ACK	acknowledgment	DSFA	data storage feature activation
AIXPCM	AIX Multipath I/O Path Control Module	DSNI	DS Network Interface
AMP	Adaptive Multi-stream Prefetching	DVE	Dynamic Volume Expansion
ANSI	American National Standards Institute	DVR	Dynamic Volume Relocation
AOS	Assist On-site	DWDM	dense wavelength division multiplexing
API	application programming interface	EAM	extent allocation method
ASIC	application-specific integrated circuit	EAV	extended address volume
ATS	Atomic Test and Set	ECC	error correction code
BF	Base Function	eDRAM	embedded dynamic random access memory
BOS	base operating system	EGP	Exterior Gateway Protocol
BPM	Backup Power Module	EKM	Encryption Key Manager
BTU	British thermal unit	EOS	end of service
CA	certificate authority	EPOW	emergency power-off warning
CCL	concurrent code load	ER	Energy Report
CDA	Code Distribution and Activation	ESCC	IBM EMEA Storage Competence Center
CEC	Central Electronics Complex	ESE	extent space efficient
CIM	Common Information Model	ESM	Enclosure Service Module
CKD	Count Key Data	FB	Fixed-Block
CMA	cable management arm	FC	Fibre Channel
CPC	central processor complex	FC-AL	Fibre Channel Arbitrated Loop
CRC	cyclic redundancy check	FCIC	Fibre Channel Interface Card
CS	Copy Services	FCP	Fibre Channel Protocol
CSI	Container Storage Interface	FDE	Full Disk Encryption
CSR	certificate signing request	FEC	Forward Error Correction
CSV	comma-separated values	FFDC	first-failure data capture
CUADD	control unit address	FICON	Fibre Channel connection
CUIR	control-unit initiated reconfiguration	FIDR	FICON Dynamic Routing
DA	device adapter	FPCCL	fast path concurrent code load
DC-UPS	direct current uninterruptible power supply	FQDN	fully qualified domain name
DDM	disk drive module	FRU	field-replaceable unit
DIF	Data Integrity Field	FSP	flexible service processor
DIMM	dual inline memory module	GbE	gigabit Ethernet
DMA	direct memory access	GC	Global Copy
DPR	Dynamic Path Reconnect	GDPS	Geographically Dispersed Parallel Sysplex
DPS	Dynamic Path Selection	GFC	gigabit Fibre Channel
DR	disaster recovery	GiB	gibibyte

GM	Global Mirror	MBps	megabytes per second
HA	high availability	MC	Management Console
HADR	high availability and disaster recovery	MDisk	managed disk
HBA	host bus adapter	ME	Management Enclosure
HCD	hardware configuration definition	MES	miscellaneous equipment specification
HDD	hard disk drive	MGM	Metro/Global Mirror
HMC	Hardware Management Console	MIB	Management Information Base
HMT	Heat Map Transfer	MiB	mebibyte
HMTU	Heat Map Transfer Utility	MIPS	million instructions per second
HPFE	High-Performance Flash Enclosure	MM	Metro Mirror
I2C	Inter-Integrated Circuit	MPIO	multipath I/O
IBM	International Business Machines Corporation	MRPD	machine-reported product data
IBM ESSNI	Enterprise Storage Server Network Interface	MSDSM	Microsoft Device Specific Module
IBM SSR	IBM Systems Service Representative	MTM	machine type and model
ICS	Install Corrective Service	MTMM	Multiple-target Metro Mirror
IML	initial microcode load	NCCL	non-concurrent code load
IOCDS	input/output configuration data set	NIST	National Institute of Standards and Technology
IODF	input/output definition file	NMP	Native Multipathing Plug-in
IOPS	I/O operations per second	NTP	Network Time Protocol
iPDU	intelligent Power Distribution Unit	NVDIMM	non-volatile dual inline memory module
IPSec	Internet Protocol Security	NVS	non-volatile storage
IRR	instruction retry and recovery	OCP	Red Hat OpenShift Container Platform
ISL	inter-switch link	ODD	OnDemand Data Dump
IWC	Intelligent Write Caching	OEL	Operating Environment License
JCL	Job Control Language	OID	object identifier
JVM	Java virtual machine	OS	operating system
KMIP	Key Management Interoperability Protocol	PAV	parallel access volume
kW	kilowatt	PCIe	Peripheral Component Interconnect Express
LAN	local area network	PCN	Power Control Network
LBA	logical block address	PFA	predictive failure analysis
LCP	Logical Corruption Protection	PFE	Product Field Engineer
LCU	logical control unit	PHYP	POWER9 IBM PowerVM Hypervisor
LD	logical device	PKCS	Public-Key Cryptography Standard
LDAP	Lightweight Directory Access Protocol	PLD	power line disturbance
LIC	Licensed Internal Code	POST	power-on self-test
LMC	Licensed Machine Code	PPRC	Peer-to-Peer Remote Copy
LPAR	logical partition	PSU	power supply unit
LSS	logical subsystem	PTC	point-in-time copy
LUN	logical unit number	PTF	Program Temporary Fix
LW	longwave	QoS	quality of service

RAID	redundant array of independent disks	SPOF	single point of failure
RAS	reliability, availability, and serviceability	SR	suspension reason
RC	reference code	SRA	Storage Replication Adapter
RCD	register clock driver	SRC	system reference code
RCL	Remote Code Load	SSD	solid-state drive
RDP	Read Diagnostic Parameters	SSH	Secure Shell
REST	Representational State Transfer	SSIC	System Storage Interoperation Center
RMZ	Remote Mirror for z/OS	SSID	subsystem identifier
RPC	Remote Procedure Call	SSL	Secure Sockets Layer
RPCC	rack power control card	SSO	single sign-on
RPM	revolutions per minute	SW	shortwave
RPQ	Request for Price Quotation	TCE	Translation Control Entry
RSC	Remote Support Center	TCO	total cost of ownership
RSCN	registered state change notification	TCT	Transparent Cloud Tiering
SAM	storage allocation method	TLS	Transport Layer Security
SAN	storage area network	TSO	Time Sharing Option
SARC	Sequential Adaptive Replacement Cache	UA	unit address
SCM	single-chip module	UCB	unit control block
SCORE	Storage Customer Opportunity Request	UPS	uninterruptible power supply
SCSI	Small Computer System Interface	URI	Unified Resource Identifier
SDD	Subsystem Device Driver	VAAI	vStorage APIs for Array Integration
SDDDSM	Subsystem Device Driver Device Specific Module	VM	virtual machine
SDDPCM	Subsystem Device Driver Path Control Module	VPD	vital product data
SDM	System Data Mover	VPN	virtual private network
SDO	Secure Data Overwrite	WLM	Workload Manager
SDS	software-defined storage	WUI	Web User Interface
SES	Small Computer System Interface Enclosure Services	WWNN	worldwide node name
SFF	small form-factor	WWPN	worldwide port name
SFI	storage facility image	zDDB	z/OS Distributed Data Backup
SFP	small form-factor pluggable	zGM	z/OS Global Mirror
SFP+	small form-factor pluggable plus	zHPF	High-Performance FICON for IBM Z
SIM	service information message	zsS	Z Synergy Services
SLB	segment lookaside buffer		
SME	subject matter expert		
SMP	symmetric multiprocessor		
SMT	simultaneous multithreading		
SNMP	Simple Network Management Protocol		
SOI	Silicon-On-Insulator		

Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide more information about the topics in this document. Some publications that are referenced in this list might be available in softcopy only.

- ▶ *Best Practices for DS8000 and z/OS HyperSwap with Copy Services Manager*, SG24-8431
- ▶ *DS8000 Global Mirror Best Practices*, REDP-5246
- ▶ *DS8870 Easy Tier Application*, REDP-5014
- ▶ *Exploring the DS8870 RESTful API Implementation*, REDP-5187
- ▶ *Getting Started with IBM zHyperLink for z/OS*, REDP-5493
- ▶ *Getting Started with IBM Z Cyber Vault*, SG24-8511
- ▶ *Getting started with z/OS Container Extensions and Docker*, SG24-8457
- ▶ *IBM Assist On-site for Storage Overview*, REDP-4889
- ▶ *IBM DS8000 Copy Services: Updated for IBM DS8000 Release 9.1*, SG24-8367
- ▶ *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-4667
- ▶ *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security (DS8000 Release 9.2)*, REDP-4500
- ▶ *IBM DS8000 High-Performance Flash Enclosure Gen2 (DS8000 R9.0)*, REDP-5422
- ▶ *IBM DS8900F and IBM Z Synergy DS8900F: Release 9.3 and z/OS 2.5*, REDP-5186
- ▶ *IBM Storage DS8000 Safeguarded Copy (Updated for DS8000 R9.2.3)*, REDP-5506
- ▶ *IBM DS8000 Transparent Cloud Tiering (DS8000 Release 9.2)*, SG24-8381
- ▶ *IBM DS8870 Easy Tier Heat Map Transfer*, REDP-5015
- ▶ *IBM DS8870 and NIST SP 800-131a Compliance*, REDP-5069
- ▶ *IBM DS8880 Thin Provisioning (Updated for Release 8.5)*, REDP-5343
- ▶ *IBM DS8900F Performance Best Practices and Monitoring*, SG24-8501
- ▶ *IBM DS8900F Product Guide Release 9.3.2*, REDP-5554
- ▶ *IBM DS8910F Model 993 Rack-Mounted Storage System Release 9.1*, REDP-5566
- ▶ *IBM Fibre Channel Endpoint Security for IBM DS8900F and IBM Z*, SG24-8455
- ▶ *IBM Power Systems S922, S914, and S924 Technical Overview and Introduction*, REDP-5497
- ▶ *IBM System Storage DS8000: Remote Pair FlashCopy (Preserve Mirror)*, REDP-4504
- ▶ *IBM System Storage DS8000: z/OS Distributed Data Backup*, REDP-4701
- ▶ *IBM z15 (8561) Technical Guide*, SG24-8851

- ▶ *IBM z15 Technical Introduction*, SG24-8850
- ▶ *IBM z16 (3931) Technical Guide*, SG24-8951
- ▶ *IBM z16 Technical Introduction*, SG24-8950
- ▶ *IBM z16 Configuration Setup*, SG24-8960
- ▶ *IBM Z Connectivity Handbook*, SG24-5444
- ▶ *LDAP Authentication for IBM Storage DS8000 Systems: Updated for DS8000 Release 9.3.2*, REDP-5460

You can search for, view, download, or order these documents and other Redbooks, Redpapers, web docs, drafts, and additional materials at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- ▶ *IBM DS8000 Host Systems Attachment Guide*, SC27-9563
- ▶ *IBM DS8000 Series Command-Line Interface User's Guide*, SC27-9562
- ▶ *IBM DS8900F Introduction and Planning Guide*, SC27-9560

Online resources

These websites are also relevant as further information sources:

- ▶ DS8000 IBM Documentation:
<https://www.ibm.com/docs/en/ds8900>
- ▶ DS8000 Series Copy Services Fibre Channel Extension Support Matrix:
<https://www.ibm.com/support/pages/ds8000-series-copy-services-fibre-channel-extension-support-matrix>
- ▶ DS8900F Code Bundle Information (includes Release Notes):
<https://www.ibm.com/support/pages/node/1072022>
- ▶ IBM DS8000 Code Recommendations:
<https://www.ibm.com/support/pages/ds8000-code-recommendation>
- ▶ IBM Fix Central:
https://www.ibm.com/support/fixcentral/swg/selectFixes?parent=Enterprise%20Storage%20Servers&product=ibm/Storage_Disk/DS8900F
- ▶ IBM Remote Code Load:
<https://www.ibm.com/support/pages/ibm-remote-code-load>
- ▶ IBM System Storage Interoperation Center (SSIC) for DS8000:
<https://www.ibm.com/systems/support/storage/ssic/>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

IBM Training and Skills

ibm.com/training

IBM Storage DS8900F Architecture and Implementation: Updated for Release 9.3.2

SG24-8456-03
ISBN 0738460753

(1.0" spine)
0.875" x 1.498"
460 x 788 pages



Redbooks



SG24-8456-03

ISBN 0738460753

Printed in U.S.A.

Get connected

