Redbooks
ibm.com/redbooks

IBM Storage DS8A00 Architecture and Implementation Guide: Updated for Release 10.1

Thorsten Altmannsberger
Sherry Brunson
Reo K B
Nielson (Nino) de Carvalho
Jeffery Cook
Tony Eriksson
Meng (Mike) Gao
Vasfi Gucer
Adam Jordan
Peter Kimmel
Flavio Gondim de Morais
Connie Riggins
Andre Luiz Coelho da Silva
Mike Stenson
Robert Tondini



Storage

IBM

Redbooks



IBM Redbooks

IBM Storage DS8A00 Architecture and Implementation Guide: Updated for Release 10.1

November 2025

SG24-8559-01

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

Second Edition (November 2025)

This edition applies to DS8000 G10 systems with Licensed Machine Code (LMC) 7.10.10 (bundle version 10.10.xxx.x), referred to as Release 10.1

© Copyright International Business Machines Corporation 2025. 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 xiii
Now you can become a published author, too! xv
Comments welcome xvi
Stay connected to IBM Redbooks xvi

Chapter 1. Introducing the IBM DS8000 G10 storage system 1
1.1 Introduction to the DS8A00 system 2
1.2 DS8A00 controller options and frames 9
1.3 DS8A00 architecture and functions overview 10
1.3.1 Overall architecture and components 10
1.3.2 Storage capacity 13
1.3.3 Supported environments 13
1.3.4 Configuration flexibility 13
1.3.5 Copy Services functions 16
1.3.6 Service and setup 16
1.3.7 IBM-Certified Secure Data Overwrite 18
1.4 Performance features 18
1.4.1 Three-layer architecture 18
1.4.2 4.8 TB cache 19
1.4.3 32 Gbps host adapters 19
1.4.4 Sophisticated caching algorithms 19
1.4.5 Flash storage 20
1.4.6 Performance for IBM Z 20
1.4.7 Performance enhancements for IBM Power servers 22

Chapter 2. IBM DS8A00 hardware components and architecture 39
2.1 Flash drive terminology of the DS8A00 40
2.1.1 Storage system 40
2.1.2 Hardware Management Console 41
2.1.3 Central processor complex 41
2.2 DS8A00 configurations and models 42
2.2.1 DS8A80 configuration 42
2.2.2 DS8A50 configuration 43
2.2.3 DS8A10 Rack configuration 46
2.2.4 DS8A10 Rack-Mounted configuration 47
2.2.5 DS8A00 base frames 49
2.2.6 DS8A00 expansion frame 52
2.2.7 Scalable upgrades 54
2.2.8 Licensed functions 54
2.3 DS8A00 architecture overview 55
2.3.1 IBM Power9 processor-based CPCs 55
2.3.2 Processor memory 57
2.3.3 Flexible service processor 58
2.3.4 Ethernet connections 58
2.3.5 Peripheral Component Interconnect Express adapters 60

© Copyright IBM Corp. 2025.

2.4 I/O enclosures and adapters 61
2.4.1 Cross-cluster communication 63
2.4.2 I/O enclosure adapters 63
2.5 Flash drive enclosures 67
2.6 Power and cooling 69
2.6.1 Intelligent Power Distribution Units 69
2.6.2 Power domains 73
2.6.3 Backup Power Modules and NVDIMM 73
2.6.4 Power cord options 74
2.6.5 Enclosure power supply units 75
2.7 Hardware Management Console and network 75

Chapter 3. Virtualization concepts 77
3.1 Virtualization definition 78
3.2 Benefits of virtualization 78
3.3 Abstraction layers for drive virtualization 79
3.3.1 Array sites 79
3.3.2 Arrays 79
3.3.3 Ranks 81
3.4 Extent pools 82
3.4.1 Dynamic extent pool merge 84
3.4.2 Logical volumes 86
3.4.3 Allocation, deletion, and modification of LUNs and CKD volumes 92
3.4.4 Volume allocation and metadata 96
3.4.5 Logical subsystems 100
3.4.6 Volume access 102
3.4.7 Virtualization hierarchy summary 104
3.5 Terminology for IBM Storage products 106

Chapter 4. IBM DS8A00 reliability, availability, and serviceability 109
4.1 DS8A00 processor complex features 110
4.1.1 Power9 PowerVM Hypervisor 110
4.1.2 Power9+ processor 110
4.1.3 Cross-cluster communication 114
4.1.4 Environmental monitoring 114
4.1.5 Resource deconfiguration 115
4.2 CPC failover and failback 116
4.2.1 Dual cluster operation and data protection 116
4.2.2 Failover 117
4.2.3 Failback 118
4.2.4 NVS and power outages 119
4.3 Data flow in the DS8A00 120
4.3.1 I/O enclosures 120
4.3.2 Host connections 120
4.3.3 Metadata checks 124
4.4 RAS on the Hardware Management Console 124
4.4.1 Licensed Internal Code updates 125
4.4.2 Call Home and remote support 125
4.5 RAS on the storage system 126
4.5.1 RAID configurations 127
4.5.2 Drive path redundancy 128
4.5.3 Flash RAID controller redundancy 128
4.5.4 Predictive failure analysis 129

4.5.5 Disk scrubbing . . . 129
4.5.6 RAID support . . . 130
4.5.7 RAID 6 overview . . . 130
4.5.8 RAID 10 overview . . . 132
4.5.9 Smart Rebuild . . . 133
4.5.10 Spare creation . . . 134
4.5.11 Services Offering with Ransomware Threat Detection . . . 135
4.6 RAS on the power subsystem . . . 135
4.6.1 Power components . . . 136
4.6.2 Line power loss . . . 137
4.6.3 Line power fluctuation . . . 137
4.6.4 Power control . . . 137
4.7 Other features . . . 138
4.7.1 Internal network . . . 138
4.7.2 Earthquake resistance . . . 138
4.7.3 IBM Certified Secure Data Overwrite . . . 139

Chapter 5. IBM DS8000 physical planning and installation . . . 143
5.1 Considerations before the installation: Planning for growth . . . 144
5.1.1 Client responsibilities for the installation . . . 144
5.1.2 Participants . . . 145
5.1.3 Required information . . . 145
5.2 Planning for the physical installation . . . 146
5.2.1 Delivery and staging area . . . 146
5.2.2 Floor type and loading . . . 147
5.2.3 Overhead cabling features . . . 148
5.2.4 Room space and service clearance . . . 149
5.2.5 Power requirements and operating environment . . . 150
5.2.6 Host interface and cables . . . 152
5.2.7 Host adapter Fibre Channel specifics for open environments . . . 155
5.2.8 FICON specifics on a z/OS environment . . . 155
5.2.9 Best practices for host adapters . . . 156
5.2.10 Worldwide node name and worldwide port name determination . . . 156
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 . . . 161
5.3.4 Remote support connection . . . 161
5.3.5 Storage area network connection . . . 161
5.3.6 Key manager servers for encryption . . . 162
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
5.6 z/OS maintenance . . . 165

Chapter 6. IBM DS8A00 Management Console planning and setup . . . 167
6.1 DS8A00 Management Console overview . . . 168
6.1.1 Management Console hardware . . . 168
6.1.2 Private and Management Ethernet networks . . . 169
6.2 Management Console software . . . 172
6.2.1 DS Storage Management GUI . . . 172

6.2.2 Data Storage Command-Line Interface 172
6.2.3 RESTful application programming interface 172
6.2.4 IBM Copy Services Manager interface 173
6.2.5 Updating the embedded IBM Copy Services Manager 173
6.3 Web User Interface 177
6.3.1 Logging in to the HMC WUI 177
6.3.2 IBM ESSNI server 179
6.4 Management Console activities 179
6.4.1 Management Console planning tasks 179
6.4.2 Planning for Licensed Internal Code upgrades 180
6.4.3 Time synchronization 181
6.4.4 Monitoring DS8A00 with the Management Console 181
6.4.5 Event notification through syslog 182
6.4.6 Call Home and remote support 182
6.5 Management Console network settings 182
6.5.1 Configuring the Management Console Network 183
6.5.2 Private networks 184
6.6 User management 184
6.6.1 Password policies 185
6.6.2 Remote authentication 186
6.6.3 Service Management Console User Management 187
6.6.4 Service Management Console LDAP authentication 192
6.6.5 Multifactor authentication (MFA) 193
6.7 Secondary Management Console 194
6.7.1 Management Console redundancy benefits 194

Chapter 7. IBM DS8000 G10 features and licensed functions 195
7.1 DS8000 licensed functions 196
7.1.1 General introduction to licensing 198
7.1.2 Licensing cost structure 200
7.2 Activating licensed functions 203
7.2.1 Obtaining DS8000 machine information and activating license keys 204
7.2.2 Obtaining the activation codes 211
7.2.3 Applying activation codes by using the DS CLI 215
7.3 Licensed scope considerations 217
7.4 Expert Care and ServicePacs 218

Chapter 8. DS8000 G10 configuration flow 221
8.1 Configuration worksheets 222
8.2 User and role management 222
8.3 Encryption 224
8.4 Network security 224
8.5 Configuration flow 225
8.6 General storage configuration guidelines 226

Chapter 9. IBM DS8000 G10 Storage Management GUI 229
9.1 Introduction 230
9.2 DS GUI: Getting started 231
9.2.1 Accessing the DS GUI 231
9.2.2 System Setup wizard 232
9.2.3 Configuring Fibre Channel port protocols and topologies 235
9.3 Managing and monitoring the storage system 237
9.3.1 Overview of DS GUI 237
9.3.2 Storage Management help functions 244

9.4 System configuration overview 245
9.4.1 Network settings 245
9.4.2 Security settings 247
9.4.3 Security settings 250
9.5 System settings 258
9.5.1 Notifications settings 262
9.5.2 Support settings 265
9.6 Logical configuration overview 268
9.7 Logical configuration for open systems volumes 269
9.7.1 Configuration flow 269
9.7.2 Creating FB pools for open systems hosts 269
9.7.3 Creating FB volumes for open systems hosts 275
9.7.4 Creating FB host attachments 279
9.7.5 Assigning FB volumes 287
9.8 Logical configuration for Count Key Data volumes 289
9.8.1 Configuration flow 289
9.8.2 Creating CKD storage pools 289
9.8.3 Creating CKD logical subsystems 292
9.8.4 Creating CKD volumes 295
9.8.5 Creating CKD parallel access volumes 296
9.8.6 Setting the FC port protocols for IBM Z attachment 300
9.9 Expanding volumes 300
9.10 Deleting a pool 302
9.11 Deleting volumes 302
9.12 Reinitializing a thin-provisioned volume 304
9.13 Easy Tier support 304
9.14 Monitoring system health 305
9.14.1 Hardware components: Status and attributes 307
9.14.2 Viewing components health and state from the system views 313
9.14.3 Monitoring system events 315
9.14.4 Exporting system-wide information 316
9.14.5 Audit logs 318
9.15 Performance monitoring 319
9.15.1 Performance statistics 320
9.15.2 Working with customized performance graphs 323
9.16 Fibre Channel error rate statistics 333
9.17 Providing feedback 334

Chapter 10. IBM DS8A00 Storage Management Command-line Interface 335
10.1 DS CLI overview 336
10.1.1 Supported operating systems for the DS CLI 336
10.1.2 Installation hints and tips 337
10.1.3 Installing the DS CLI on a Windows 11 system 338
10.1.4 Installing the DS CLI on an z/OS system 338
10.1.5 DS CLI version 343
10.1.6 User accounts 344
10.1.7 User management by using the DS CLI 344
10.1.8 DS CLI profile 346
10.1.9 Configuring the DS CLI to use the second HMC 349
10.1.10 Command structure 349
10.1.11 Using the DS CLI application 350
10.1.12 Return codes 352
10.1.13 User assistance 352

10.2 I/O port configuration . . . 353
10.3 DS8A00 storage configuration for Fixed-Block volumes . . . 354
10.3.1 Disk classes . . . 354
10.3.2 Creating the arrays . . . 354
10.3.3 Creating the ranks . . . 356
10.3.4 Creating the extent pools and assigning ranks to the extent pools . . . 357
10.3.5 Creating the FB volumes . . . 359
10.3.6 Creating the volume groups . . . 366
10.3.7 Creating host connections and clusters . . . 368
10.3.8 Mapping open system host disks to storage unit volumes . . . 375
10.4 DS8000 storage configuration for the CKD volumes . . . 375
10.4.1 Disk classes . . . 376
10.4.2 Creating the arrays . . . 376
10.4.3 Creating the ranks and extent pools . . . 376
10.4.4 Creating the extent pools . . . 377
10.4.5 Logical control unit creation . . . 379
10.4.6 Creating the CKD volumes . . . 380
10.4.7 Resource groups . . . 388
10.4.8 IBM Easy Tier . . . 388
10.5 Metrics with DS CLI . . . 388
10.5.1 Overview of metrics . . . 389
10.5.2 Offloading performance data and other parameters . . . 395
10.6 Private network security commands . . . 396
10.7 Copy Services commands . . . 398
10.8 Earlier DS CLI commands and scripts . . . 398
10.9 Dual Control commands . . . 399
10.9.1 Examples of Dual Control DSCLI commands . . . 399
10.10 For more information . . . 402

Chapter 11. Licensed Machine Code . . . 405
11.1 How new Licensed Internal Code is released . . . 406
11.2 Bundle installation . . . 408
11.2.1 Code update methods . . . 409
11.2.2 Remote Code Load . . . 412
11.2.3 Customer Code Load . . . 414
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 . . . 420
11.6 Postinstallation activities . . . 421
11.7 Summary . . . 421

Chapter 12. Monitoring and support . . . 423
12.1 SNMP implementation on the DS8000 . . . 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
12.2.4 Compression SNMP . . . 432

12.3 SNMP configuration 433
12.3.1 SNMP preparation 433
12.3.2 SNMP configuration with the HMC 433
12.3.3 SNMP configuration with the DS CLI 437
12.4 Introducing remote support 437
12.5 IBM policies for remote support 438
12.6 Remote support advantages 438
12.7 Remote support and Call Home 438
12.7.1 Call Home and heartbeat: Outbound 438
12.7.2 Data offload: Outbound 439
12.7.3 Outbound connection types 439
12.8 Remote Support Access (inbound) 440
12.8.1 Assist On-site 440
12.8.2 DS8A00-embedded AOS 441
12.8.3 IBM Remote Support Center for DS8000 442
12.8.4 Support access management through the DS CLI and DS GUI 442
12.9 Call Home and Assist On-Site customer-provided certificates 445
12.10 RESTful API 447
12.11 Audit logging 450
12.12 Using IBM Storage Insights 451
12.12.1 IBM Storage Insights 452
12.12.2 Getting started with IBM Storage Insights 452
12.12.3 Case interaction with IBM Storage Insights 455
12.12.4 IBM Storage Insights: Alert policies 457
12.13 IBM Call Home Connect Cloud 459

Chapter 13. Optimizing performance: Practical tips and references 461
13.1 Data flow management and control 462
13.1.1 Choosing hardware 462
13.1.2 Workload allocation to hardware 464
13.2 Planning and analyzing arrays and RAID 464
13.2.1 Logical configuration overview 464
13.2.2 Logical configuration performance considerations 465
13.3 Understanding your workload 467
13.3.1 General workload types 467
13.3.2 Database workload 469
13.3.3 Application workload 470
13.3.4 Profiling workloads in the design phase 471
13.3.5 Workload isolation versus resource sharing 471
13.4 Performance planning tool: IBM Storage Modeller 473
13.5 Hardware planning for performance 473
13.5.1 Host Adapters 473
13.5.2 Device Adapters 473
13.5.3 High-Performance Flash Enclosure - Gen3 473
13.5.4 zHyperLink 474
13.5.5 Multipathing 474
13.5.6 Host attachment spreading 474
13.5.7 Caching 475
13.6 IBM Z host performance 475
13.6.1 Performance monitoring with RMF 475
13.6.2 Z Systems planning and configuration 476
13.6.3 Optimizing performance 477
13.7 IBM i host performance 478

13.7.1 IBM i storage architecture 478
13.7.2 Fibre channel adapters and multipath 478
13.7.3 Performance guidelines for drives 479
13.7.4 Number and size of LUNs 479
13.7.5 Analyzing performance data 480
13.8 Host adapters 481
13.8.1 Fibre channel-attached hosts and SAN 481
13.8.2 FICON and IBM Z hosts 483
13.9 Practical performance management 485

Chapter 14. Ransomware Threat Detection 487
14.1 Ransomware Threat Detection (RTD) overview 488
14.1.1 Key benefits 488
14.1.2 Example scenario 489
14.2 RTD logical model 489
14.3 Configure Threat Detection Storage (TDS) 491
14.3.1 Security Group License for IBM DS8000 systems 491
14.3.2 Automatic Sending Threat Detection Data to IBM 491
14.3.3 Call Home settings 492
14.3.4 Threat Detection Storage reference 493
14.4 Storage Insights Pro for alerting 494
14.4.1 Ransomware alert details 494
14.5 Comparing IBM FlashSystem and DS8000: Key differences in Ransomware Threat Detection 498

Abbreviations and acronyms 501

Related publications 505
IBM Redbooks 505
Other publications 506
Online resources 506
Help from IBM 507

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.

© Copyright IBM Corp. 2025.

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 https://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 FlashCore®	PowerVM®
DS8000®	IBM FlashSystem®	Q®
Easy Tier®	IBM Research®	RACF®
EAV®	IBM Security®	Redbooks®
FICON®	IBM Services®	Redbooks (logo) ®
FlashCopy®	IBM Spectrum®	ServicePac®
GDPS®	IBM Sterling®	Sterling™
Guardium®	IBM Z®	z/OS®
HyperSwap®	Parallel Sysplex®	z/VM®
IBM®	Power Architecture®	z16®
IBM Cloud®	Power8®	z17™
IBM Cloud Pak®	Power9®	zSystems®

The following terms are trademarks of other companies:

ITIL is a Registered Trade Mark of AXELOS Limited.

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 DS8A00 family. The book provides reference information to assist readers who need to plan for, install, and configure the IBM DS8000® G10 systems. This edition applies to DS8000 G10 systems with Licensed Machine Code (LMC) 7.10.10 (bundle version 10.10.*xxx.x*), referred to as Release 10.1.

The DS8000 G10 systems are all-flash exclusively, and these three classes and four models are offered:

- IBM DS8A80: Analytic Class

  The DS8A80 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 DS8A50: Agility Class

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

- IBM DS8A10: Flexibility Class

  The Flexibility Class reduces complexity and addresses various workloads at the lowest DS8000 family entry cost.
  The DS8A00 is available as either a racked or rack-mountable model.

The DS8A00 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) Gen-3.

Like its predecessors, the DS8A00 models support advanced disaster recovery (DR) solutions, business continuity solutions, and thin provisioning. In addition to several performance enhancements, they offer compression and real-time threat detection.

The predecessor DS8900F models are described in *IBM Storage DS8900F Architecture and Implementation*, SG24-8456, and *IBM DS8910F Model 993 Rack-Mounted Storage System*, REDP-5566.

Authors

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

Thorsten Altmannsberger is a Senior System Engineer specializing in IBM storage solutions, with a strong focus on mainframe storage technologies, including both disk and tape systems. Since 2023, he has been working for SVA, an IBM Business Partner based in Wiesbaden, Germany. From 1996 to 2023, he held various roles at IBM - starting as a Product Field Engineer providing last-level tape support, and since 2005 serving as a Technical Specialist in storage presales. He earned a degree in Electrical Engineering from the University of Applied Sciences in Frankfurt am Main, Germany.

© Copyright IBM Corp. 2025.

Sherry Brunson joined IBM in March of 1985 and worked as a large system IBM service representative before becoming a Top Gun in 1990. Sherry is a Top Gun in the Eastern US for all storage products, Power systems, and IBM Z Systems. She has supported and implemented DS8000 and Scaled out Network Appliance storage products globally as well as developing and teaching educational classes. She has also taught Z System classes in the United States.

Reo K B is a Delivery Consultant at IBM India Expert Labs, specializing in DS8000 and FlashSystem storage solutions. He joined IBM in 2019 and brings over 17 years of experience in storage systems. Reo holds a Bachelor's degree in Electrical and Electronics Engineering from Anna University, India**.**

Nielson (Nino) de Carvalho is an IBM Certified Specialist at Expert level based out of Switzerland, with over a decade of experience in IBM Mainframe computing. He specializes in IBM Z, LinuxONE, z/OS®, z/VM®, GDPS®, DS8000, TS7700, and SAN solutions. Nino has extensive technical experience serving IBM clients across various sectors globally and has co-authored several IBM Redbook publications.

Jeff Cook is a DS8000 Subject Matter Expert (SME), and leads the Tucson, Arizona DS8000 Product Engineering team. He has been with IBM for 46 years, providing implementation and technical support to customers and service representatives in complex enterprise environments. For the past 27 years, he has specialized in IBM direct access storage device products, specifically the DS8000 and former enterprise storage systems.

Tony Eriksson works at IBM Systems Lab Services Nordic, based in Stockholm, Sweden. He joined IBM in 1988 as a mainframe Customer Engineer. From 1999, he has worked as a product specialist on enterprise storage solutions, including IBM DS8000, IBM SAN Volume Controller, and SAN products. In 2014, he joined an IBM Business Partner as a pre-sales and implementation consultant, mainly migrating to IBM storage products. In 2019, he joined IBM Systems Lab Services Nordic working close with the Swedish technical sales team for storage solutions. He holds a Technical College Graduation in Electrical Engineering.

Meng (Mike) Gao is a DS8000 Subject Matter Expert and has been a member of the DS8000 remote support team since 2020. Previously, he was part of the System Services Representative team, supporting all IBM logo products as well as MVS products since 2013. He has extensive experience working closely with clients to troubleshoot hardware and software issues related to DS8000, including Copy Services and Copy Services Manager. Prior to joining IBM, he was a Data Center Engineer for one of the major banks in the U.S. He has over 15 years of experience with IT infrastructures and holds degrees in Computer and Network Security (B.S.) and Information Systems Technology (M.S.).

Vasfi Gucer leads projects for the IBM Redbooks team, leveraging his 20+ years of experience in systems management, networking, and software. A prolific writer and global IBM instructor, his focus has shifted to storage and cloud computing in the past 10 years. Vasfi holds multiple certifications, including IBM Certified Senior IT Specialist, PMP, ITIL V2 Manager, and ITIL V3 Expert.

Adam Jordan began his IBM career in 2013 as an SSR in Boulder, Colorado. He initially worked with IBM System x before transitioning to DS8000 and TS7700 service. In 2023, he joined the National Top Gun team, specializing in IBM Z, DS8000, and TS7700 systems. Adam also leads the IBM Z Hands-On seminar in Tucson, Arizona.

Peter Kimmel is a Senior Platform Engineer for Enterprise Storage in the IBM EMEA Client Engineering team 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 has co-authored several DS8000 IBM publications. He holds a Diploma (MSc) degree in physics from the University of Kaiserslautern.

Flavio Gondim de Morais is a Senior L3 Engineer involved in DS8000 development. He previously worked as a Lab Services Consultant for IBM Brazil. His expertise focuses on performance and Copy Services, and he has co-authored earlier versions of DS8000 publications. Flavio holds a Bachelor's degree in Computer Engineering and joined IBM in 2009.

Connie Riggins is a DS8000 Copy Services and Copy Services Manager Subject Matter Expert with the DS8000 Product Engineering group. She started working at IBM in 2015. Prior to joining IBM, starting in 1991, Connie worked at Amdahl Corp. as a Systems Engineer and later at Softek Storage Solutions as Product Manager of TDMF for z/OS.

Andre Coelho (Andre Luiz Coelho da Silva) is an IBM Certified Specialist Level 3 based in Brazil. He brings over 39 years of experience with z/OS and its predecessors. His areas of expertise include Data Facility Storage Management Subsystem (DFSMS), DFSMShsm (Hierarchical Storage Manager), DFSMSrmm (Removable Media Management), DS8000, and ATL/VTS. Andre has contributed to numerous IBM Redbooks publications throughout his career.

Mike Stenson is the team lead for the DS8000 development field screen team. He has over 20 years of experience of engineering, and development in storage, server, and networking environments. He joined IBM in 2005 and has worked on every generation of the DS8000. He has co-authored several IBM publications.

Robert Tondini is a mainframe storage Subject Matter Expert based in IBM Australia and New Zealand. He has 29 years of experience in mainframe storage implementation, management, capacity and performance sizing with focus on designing and exploiting high availability, disaster recovery and cyber resiliency solutions. He co-authored several IBM Redbooks publications and workshops for DS8000 systems.

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

Brian Rinaldi, John Bernatz, Stephanie Beddow, Randy Blea, Jed Dyreng, David Valverde, Matt Borlick
IBM USA

Laura Saldana
IBM Mexico

Sumesh Pahula, Ramakrishna Vadla, Anil Shetty, Vijay Patidar, Girish Baxi
IBM India

Carsten Haag
IBM Germany

Radoslav Neshev
IBM Bulgaria

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:

ibm.com/redbooks/residencies.html

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:

  ibm.com/redbooks

- 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:

  https://www.linkedin.com/groups/2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/subscribe

- Stay current on recent Redbooks publications with RSS Feeds:

  https://www.redbooks.ibm.com/rss.html

1

Introducing the IBM DS8000 G10 storage system

This chapter introduces the features and functions of the IBM DS8000 G10 generation of storage systems and its DS8A00 models that are available with IBM DS8000 Licensed Machine Code (LMC) 7.10.10 (bundle version 10.10.*xxx.x*), which is referred to as Release 10.1. More information about functions and features is provided in subsequent chapters.

This chapter covers the following topics:

- “Introduction to the DS8A00 system” on page 2
- “DS8A00 controller options and frames” on page 9
- “DS8A00 architecture and functions overview” on page 10
- “Performance features” on page 18

© Copyright IBM Corp. 2025. All rights reserved.

1.1 Introduction to the DS8A00 system

The IBM DS8000 G10 tenth generation is a high-performance, high-capacity, highly secure, and resilient series of DS8000 storage systems. Its DS8A00 model family is the most advanced of the DS8000 offerings to date. All DS8A00 models are compression-capable, all-flash systems. The high-availability (HA) multiplatform support, including IBM Z, and the simplified management tools provide a cost-effective path to an on-demand world.

Figure 1-1 shows a single-frame DS8000 G10.



Figure 1-1 Single-frame DS8000 G10

All models of this DS8000 generation belong to the IBM DS8A00 high-performance model family:

- IBM DS8A80: Analytic Class
- IBM DS8A50: Agility Class
- IBM DS8A10: Flexibility Class

The wide range of models meets the needs of any enterprise storage client, while latest-generation technology leads the transition of enterprise data storage to the new era of cyber resilience and AI. The DS8A00 models support the most demanding business

applications with their exceptional all-around performance and data throughput. Some models are shown in Figure 1-2.



Figure 1-2 DS8A00 frames (right: Tiled standard doors 2-frame, with side covers; left: Space-saver door)

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

- High performance
- High capacity
- High Availability
- Security and encryption
- Cost efficiency
- Energy efficiency
- Scalability
- Business continuity and data-protection functions.

The DS8000 G10 architecture is server-based. Two 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), which also come in a new generation.

The DS8A00 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. It can be equipped with compression-capable drives, which reduce this footprint even further.

The DS8000 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 DS8A00 to achieve the highest energy efficiency in the DS8000 series. The DS8A00 comes with Titanium power supplies which comply with latest regulations for an enhanced energy efficiency. All models are available for both three-phase and single-phase power attachment.

Figure 1-3 shows the various components within the base frame. The expansion frame which is an option for the bigger model enables clients to add more capacity to their storage systems in once again the same footprint.

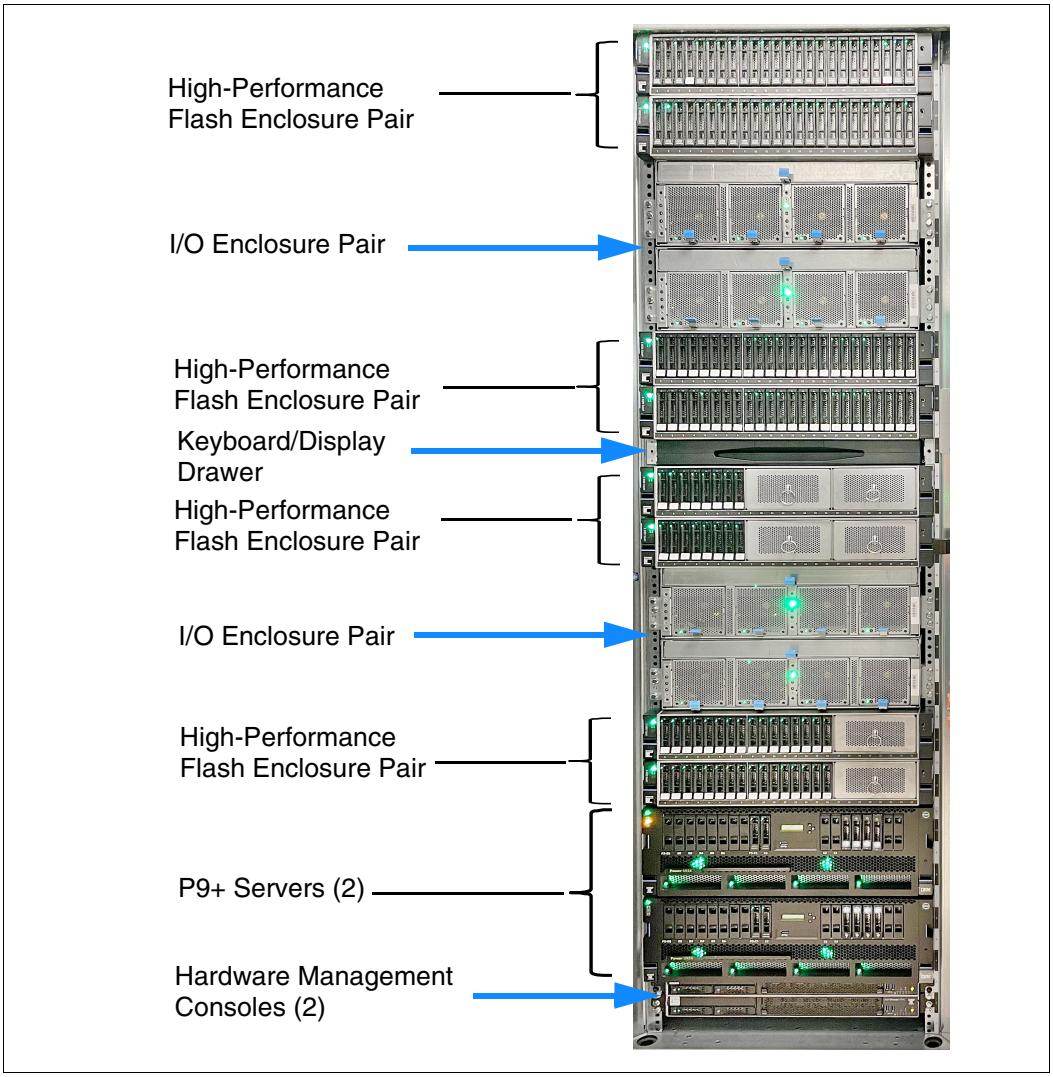


Figure 1-3 DS8A50 components (base frame)

I/O enclosures are attached to the Power9+ processor-based servers with Peripheral Component Interconnect Express (PCIe) Generation-4 cables. The I/O enclosure has six PCIe adapter slots for host/device adapters, and two slots for zHyperLink ports.

These racked models also have an integrated keyboard and display that can be accessed from the front of the rack. A pair of 1U Hardware Management Consoles (HMCs) are installed in the base rack management enclosure. The height of the rack is 42U for all units.

At-a-glance features of the DS8A00

The DS8A00 offers the following features (more details can be found in subsequent chapters of this book):

- The DS8A00 is available with various Power9+ processor options, starting with the smaller DS8A10 model, through the DS8A50, and then up to the largest model DS8A80 with its 48 cores. Each system covers a wide range of performance and cost options.

- By using the IBM Power Simultaneous Multithreading (SMT) technology, analytic data processing clients can achieve a few 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), all while minimizing latency.
- Memory configurations between 256 GB and 4.8 TB are available.
- The DS8000 is designed for up to 99.999999% (“eight nines”) availability and a no-data-loss automatic failover within metropolitan distances, for clients using best practices and an enabled Parallel Sysplex®, GDPS (Geographically Dispersed Parallel Sysplex), HyperSwap® and System Recovery Boost (SRB).
- The DS8A00 uses High-Performance Flash Enclosures: HPFE Gen3. One HPFE pair can hold 16, 32, or 48 flash drives, with nominal capacities between 1.6 TB and 19.2 TB. Some of the drives are industry-standard flash drives. The larger capacities come as IBM FlashCore Modules (FCM) which offer hardware compression. At the time of writing, the FCMs are in their 4th generation.
- Ransomware Threat Detection, integrated with embedded AI-driven FCM4 drives, can now leverage a brand-new service to spot the most advanced ransomwares in real-time. With built-in machine learning models, organizations can start collecting and analyzing data patterns and statistics to further enhance their storage defense. See Chapter 14, “Ransomware Threat Detection” on page 487 for more information.
- With long-distance replication, an RPO of 2 seconds is possible. In replication environments configured with GDPS Continuous Availability, the DS8000 supports an RTO of less than 60 seconds at distances of more than 1,000 miles.
- IBM Z Synergy Services (zsS) include the following IBM z/OS licensed features that are supported on the storage system:
  - Fibre Channel connection (IBM FICON®) attachment for IBM Z
  - Parallel access volume (PAV), HyperPAV, and SuperPAV
  - High-Performance FICON for IBM Z (zHPF)
  - zHyperLink and zHyperWrite usage
  - z/OS Distributed Data Backup (zDDB)
  - Encryption of data in flight
  - Compression with the FlashCore Modules
  - An option for thin-provisioned Count Key Data (CKD) volumes
- As another option, Transparent Cloud Tiering (TCT) is available for these clients. TCT enables hybrid cloud storage solutions as additional tiers for disaster recovery, data archiving, and long-term retention in IBM Z environments. By using TCT, you can configure direct data movement from the DS8000 CPCs to tape, to other local storage like IBM Cloud® Object Storage, or to external cloud storage, which helps reduce backup workload on IBM Z. Data can be encrypted before it leaves the DS8000. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

  TCT offers 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. If configured as object storage, the TS7700 can still be used with traditional FICON logical tape volumes. For more information, see *IBM DS8000 Transparent Cloud Tiering*, SG24-8381.
- DS8000 G10 offers enhanced connectivity with 4-port 32 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. Further 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 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).

- The DS8000 G10 includes PCIe-based I/O enclosures. I/O enclosures come as pairs. Each single I/O enclosure connects to both internal servers through a pair of x8 PCIe Gen4 cables. Each cable provides 16 GBps connectivity. Each I/O enclosure also provides two zHyperLink ports.

  Each pair of I/O enclosures supports up to two pairs of HPFEs. You can use the storage virtualization that is offered by the DS8000 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.
- The DS8000 G10 provides a graphical management interface to configure, query status information, and monitor the performance of the DS8000. The DS8000 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 previously 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 DS8000 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) helps simplify 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 DS8000 Thin Provisioning*, 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.

- An alternative Local Key Management for FDE is also available on the storage system, and it does not require any external key manager server or extra license. Here, the internal keys are stored in the encrypted Platform Key Store (PKS) areas of the Power9+ CPC server LPARs.
- The DS8000 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 any DS8A00.
- All drives of this DS8000 generation are capable of complying with FIPS 140-3, which is the latest and updated standard for security requirements for cryptographic modules, from the US National Institute of Standards and Technology.
- Payment Card Industry Data Security Standard (PCI DSS) is a feature of the encryption key management process that helps address payment industry requirements.
- DS8000 host adapters include support for IBM Fibre Channel Endpoint Security (IFCES) as part of the cybersecurity solutions that are offered by IBM. Traffic to the host can be authenticated and encrypted. As with data-at-rest encryption, the IFCES data-in-flight encryption comes without any impact to performance. To take advantage of this capability, the platforms that are attached to the DS8000 must also support IBM Fibre Channel Endpoint Security. IBM z17™ is an example. IBM Fibre Channel Endpoint Security is described in *IBM Fibre Channel Endpoint Security for IBM DS8900F and IBM Z*, SG24-8455.
- 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. The DS8000 uses a centralized LDAP directory rather than a local user repository. The DS8A00 supports native LDAP configuration, and it does not require any IBM CSM proxy. 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 DS8000 provides 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.
- For customers using Copy Services (CS feature), the IBM Copy Services Manager (CSM) software is freely available to install and use on the Hardware Management Consoles. CSM helps to manage complex replication environments, like multi-site PPRC and their combinations with IBM FlashCopy® or Safeguarded Copy.
- The IBM Safeguarded Copy functionality protects data from malware, ransomware, and many other cyberattacks. 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*, 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 DS8000 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 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.
- Features like Dual Control for destructive actions, Secure Boot, or z/OS PassTicket support for password-less authentication give extra security and peace of mind.
- The DS8000 makes use of synergies with IBM Z in more than 30 advanced capabilities and includes the following DS8000 G10 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 IBM zSystems® 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 Z supports SAN authentication and FC encryption of FICON traffic to the DS8000.
  - IBM zHyperLink

    zHyperLink is a short-distance link technology that dramatically reduces latency by connecting the IBM Z CPCs directly to the I/O bays of the DS8000.

    All DS8A00 models support zHyperLink, for both read and write I/O. zHyperLink is intended to complement FICON technology to accelerate I/O requests that are typically used for transaction processing.

    The DS8A00 models accelerate zHyperLink one step further compared to previous DS8000 generations: For instance, the transfer sizes for reads of qualified I/Os can now grow up to 32K, which allows further usage scenarios and enables higher throughputs.

    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 DS8000 and IBM Z Synergy*, REDP-5186.

1.2 DS8A00 controller options and frames

The DS8A00 family consists of distinct DS8000 G10 models of all-flash storage:

- DS8A10 machine type 5341 model A00 Rack-Mount, which is to be installed in a client-provided standard rack
- DS8A10 machine type 5341 model A01 as racked in a single frame
- DS8A50 machine type 5341 model A05 with a model E05 expansion frame option
- DS8A80 machine type 5341 model A08 with a model E05 expansion frame option

The DS8A00 includes the following features:

- IBM Power9+ processor-based server technology

  The DS8A00 features the IBM Power9+ processor-based server technology for high performance. Compared to the IBM Power8® processors that were used in the DS8880, IBM Power9®, and 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 is almost doubled between these two Power generations.

- PCIe Gen-4

  With the change of the servers from Power9 to Power9+, the internal connectivity within the DS8A00 is upgraded to PCIe Gen-4, which typically doubles throughput of the respective connection.

- Nondisruptive upgrade path

  Upgrades to the DS8000 are non-disruptive. You can upgrade processor, cache, host adapters, storage, and expansion frames concurrently without disrupting applications.

- The airflow system ensures optimal horizontal cooling of the storage system. The DS8000 is designed for hot-aisle and cold-aisle data center configurations, drawing cool air from the front and exhausting warm air at the rear. For more information, see 2.6, “Power and cooling” on page 69.

- The DS8A00 offers three system-class options:
  - DS8A80 Model A08 – Analytic Class, which is based on the IBM Power S924 architecture.
  - DS8A50 Model A05 – Agility Class, which again is based on the IBM Power S924 architecture.
  - DS8A10 Models A00 and A01– 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 55.

- HPFE

  HPFE Gen3 is an enclosure that can support high-performance and high-capacity flash drives that can provide a capacity of 1.6 TB–19.2 TB (2.5-inch form-factor) in a 2U rack space. The enclosures must be installed in pairs. All drive media have NVMe interfaces, and some of them are compression-capable.

- Rack-mountable and racked options

  The DS8A00 comes in racks with a footprint that is slightly less than a DS8900F and much smaller than its DS8880 predecessor models. To use a DS8000 in a customer-provided rack, order the rack-mount DS8A10 A00 model.

1.3 DS8A00 architecture and functions overview

The DS8000 G10 and its DS8A00 models offer continuity with the fundamental architecture of its predecessors: the DS8900F and DS8880 models. This architecture ensures that the DS8A00 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 DS8A00, see Chapter 2, “IBM DS8A00 hardware components and architecture” on page 39.

IBM POWER processor technology

The Power9+ and Power9 processors are manufactured by using the 14 nm Silicon-On-Insulator (SOI) technology. Each chip is 695 mm^2 and contains 8 billion transistors. The DS8A00 uses the PCIe I/O controllers and an interconnection system that connects all components within the chip. The POWER processor advancements in multi-core and multithreading are significant. These multithreading capabilities improve the I/O throughput of the DS8000 storage servers.

The DS8A00 family offers several configurations of CPCs:

- The DS8A80 offers a dual-24-core configuration (48 cores total) with 4.8 TB of total system memory
- The DS8A50 offers dual-20-core configurations (40 cores total) with a total 1 TB, 2 TB, or 3.5 TB of system memory
- The DS8A10 has a 10-core processor configuration per CPC (20 cores total) with 256 GB or 512 GB of total system memory.

A CPC is also referred to as a *storage server*. For more information, see Chapter 4, “IBM DS8A00 reliability, availability, and serviceability” on page 109.

Internal PCIe-based fabric

The DS8000 fabric includes the following specifications:

- DS8A00 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 Gen4 connectivity to all installed host and device adapters. Each I/O enclosure features six PCIe Gen4 x16 adapter slots for these, two for the device adapters and two connectors for attachment to zHyperLink optical transceiver ports, and another four PCIe x8 connecting to host adapters.

For more information, see Chapter 2, “IBM DS8A00 hardware components and architecture” on page 39.

HPFE Gen3

The HPFE flash RAID adapters are installed in pairs and split across an I/O enclosure pair. They occupy the respective top 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 PCIe Gen4 cables for high bandwidth and redundancy. Each drive enclosure can contain up to twenty-four 2.5-inch NVMe flash drives or modules. The flash devices are installed in groups of 16, and split evenly across the two drive enclosures in the pair.

Each flash adapter pair and HPFE pair delivers a throughput in a range of around 20 GBps, and for small blocks, with reads, more than a million IOPS.

For more information, see Chapter 2, “IBM DS8A00 hardware components and architecture” on page 39.

Drive options

Flash drives use less power than HDDs, and their IOPS count per drive can be up to 100 times higher with a 10 times shorter response time. Drive vendors also offer flash drives with various specifications, such as enterprise versus consumer, single-level versus multi-level cell, endurance, and DRAM cache, and each generation is a little more advanced.

Flash drives are grouped into three classes, based on performance, capacity, and compression capability. The following 2.5-inch NVMe flash devices are currently supported with DS8A00:

- Compression Tier – IBM FlashCore® Modules:
  - 4.8 TB
  - 9.6 TB
  - 19.2 TB
- Capacity Tier 1 – Industry-standard flash drives:
  - 1.6 TB
- Capacity Tier 2 – Industry-standard flash drives:
  - 1.92 TB
  - 3.84 TB
  - 7.68 TB

All flash drives in the DS8000 Gen10 are encryption-capable and always encrypt data internally. Full encryption enablement is optional and requires either local key management or at least two external key servers.

IBM FlashCore Modules

The IBM FlashCore Modules (FCM) feature hardware-enabled compression and encryption. IBM already offers them in their fourth generation. With an internal mix and tiering of Single-Level Cell (SLC) and Quad-Level Cell (QLC) flash they deliver a higher density and capacity and maintain the same excellent response times. FCMs are computational storage and come with a specialized processor, associated memory and a highly-adapted firmware, which is not available in industry-standard flash. In addition to compression, their extra compute power also allows AI-enabled real-time inline ransomware threat detection.Sudden, unusual changes in access patterns or data structures can be detected quickly and may indicate a potential attack.

For more information on FCM Gen-4, see *IBM FlashCore Module (FCM) Product Guide*, REDP-5725.

Easy Tier

The DS8000 can use Easy Tier to automatically balance data placement on disk drives to avoid hot spots on flash arrays. Easy Tier can place data on a storage array that is less loaded and best suits the access frequency of the data. For certain drive combinations, such as between the capacity tiers of the industry-standard drives, data can be moved nondisruptively to a higher or lower tier.

For more information about Easy Tier, see *IBM DS8000 Easy Tier*, REDP-4667.

Host adapters

The DS8A00 offers 32 Gbps host adapters. They have four ports each, and each port can be independently configured for either FCP or FICON. Each port independently auto-negotiates to a 32 Gbps, 16 Gbps, or 8 Gbps link speed.

For more information, see Chapter 2, “IBM DS8A00 hardware components and architecture” on page 39.

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 DS8000 includes two HMCs for redundancy, which are installed at the bottom in the base rack. DS8A00 HMCs support IPv4 and IPv6 standards. For more information, see Chapter 6, “IBM DS8A00 Management Console planning and setup” on page 167.

Isolated key server

The DS8000 includes FDE flash drives which internally always encrypt. To configure a DS8000 to use data-at-rest encryption, either Local Encryption, or at least two external key servers are required. In the latter case, at least one isolated key server with dedicated hardware and non-encrypted storage resources is required. For more information, see 5.3.6, “Key manager servers for encryption” on page 162.

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

Local Key Manager

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 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*, REDP-4500.

1.3.2 Storage capacity

The physical storage capacity for the DS8A00 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 FlashCore Modules and industry-standard 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 or RAID 10. RAID 6 is the default and preferred setting for the DS8A00. As of the time of this writing, a drive extent pool consists of either FCM only, or industry-standard drives only.

For more information, see 2.2, “DS8A00 configurations and models” on page 42.

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 DS8A00 supports over 60 platforms. For the list of supported platforms, see the IBM System Storage Interoperation Center (SSIC) for IBM Storage Enterprise Disk/DS8A00.

This support of heterogeneous environments and attachments, with the flexibility to partition 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 3, “Virtualization concepts” on page 77.

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 DS8000 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 to improve 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 DS8A00 supports features in the OpenStack environment, such as volume replication and volume retype. The Cinder driver for DS8000 is 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.

IBM DS8000 provides persistent storage for mission-critical containers and supports IBM Cloud Pak® solutions, enhancing and extending the capabilities of Red Hat OpenShift.

The DS8000 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 block storage CSI driver documentation.

RESTful application programming interface

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

For more information about the RESTful API, see the *RESTful API Guide*, SC27-9823, or *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 to better use the 256 address range, particularly for IBM Z.

Simplified LUN masking

In the GUI, LUNs are directly mapped to the host, and the user cannot define volume groups. Volume groups still exist on the DS8000 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, and compression

Volumes in the DS8000 G10 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.

Thin-provisioned volumes provide over-provisioning capabilities that provide more efficient usage of the storage capacity and reduced storage management requirements. For more information, see Chapter 3, “Virtualization concepts” on page 77 and *IBM DS8000 Thin Provisioning*, REDP-5343.

Data in the FCM Compression Tier is always stored thin and compressed. New or enhanced system commands, GUI views, warning thresholds, and alerts are available to monitor and manage the available capacities. When you use over-provisioning, define a policy to handle these alerts and let IBM help you with an appropriate planning for the right capacities and procedures.

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. The DS8000's IBM FlashCopy technology provides near-instantaneous point-in-time copies (PTCs) of data, ideal for IT environments that cannot tolerate system downtime for traditional backups. Also, The DS8000 can 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, which is based on the FlashCopy technology, can create and retain hundreds of PTCs for protection against logical data corruption. In PPRC-replicated DS8000 pairs, you can also incrementally restore a recovered Safeguarded Copy to a production copy of the data. With the additional integration level to a full Cyber Vault solution you could automate and accelerate detection and response to a cyberattack. For more information, see *IBM Storage DS8000 Safeguarded Copy*, REDP-5506, and *Getting Started with IBM Z Cyber Vault,* SG24-8511. Since Release 10.1, the maximum number of Safeguarded Copy backups that can be created has been increased to 1024 copies.

For data protection and availability needs, the DS8000 provides MM, GM, GC, and MGM, 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*, SG24-8367.

1.3.6 Service and setup

DS8000 systems have a single machine type that is called 5341. Warranty and service options are offered as part of Expert Care. Options range from a 1-year Expert Care Advanced selection up to a 5-year Expert Care Premium, and can be chosen when you order a new system or an upgrade as requirements change.

Figure 1-4 describes the available Expert Care options.

Expert Care Services	Base Warranty 5341	Advanced 5130 (A01–A05)	Premium 5130 (P01–P05)
Hardware Maintenance (IOR = IBM On-site Repair)	1-Year 24×7 same day	24×7 same day	24×7 same day
Support Line (24×7 remote technical support)		✓	✓
Predictive Support		✓ SI	✓ SI Pro
Technical Account Manager (TAM)			✓
Enhanced Response Time (30 min for Severity 1 & 2)			✓
Remote Code Load (up to 2× per year)		Optional add-on	✓
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 to 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 each DS8000 Base rack, 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.

You can also select IBM Remote Support Center (RSC) so that IBM Support can access your systems remotely.

IBM provides three options for microcode updates:

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

For customers who choose Expert Care Advanced or the base warranty service, Customer-controlled 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 can 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.

Customers who want to have an IBM Systems Service Representative (IBM SSR) perform Onsite Code Load can purchase Feature Code #AHY2 with Expert Care Premium, or Feature Code #AHY3 with Expert Care Advanced. Customers selecting EC Advanced but still opting for RCL can select Feature Code #AHY4.

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 DS8000 storage system. Before you perform an SDO, remove all the logical configuration and any existing encryption groups. The process is then initiated by the IBM SSR. For more information, see 4.7.3, “IBM Certified Secure Data Overwrite” on page 139.

1.4 Performance features

The DS8000 offers optimally balanced performance. The DS8A00 incorporates many performance enhancements, such as a dual multi-core IBM Power9+ processor complex implementation, fast 32 Gbps FCP / FICON host adapters, HPFE dedicated flash architecture in third generation, classical flash drives or hardware-compression enabled flash modules, and high-bandwidth, fault-tolerant point-to-point PCIe Gen-4 internal connections, using NVMe.

With all of these components, the DS8A00 is positioned at the top of the high-performance category.

1.4.1 Three-layer architecture

In a new 3-layer, shared-everything architecture, the DS8000 G10 not only has a strong core with its central processor complexes (CPC), but also strong peripheral layers, such as the host adapter or device adapter levels or with the computational storage in the FlashCore Modules. These each have very strong and independent computational capabilities.

This approach not only relieves the CPC units from additional work, but more importantly, it minimizes unnecessary data transfers by performing computations close to where they are needed. As a result, internal bandwidth is conserved, and the offloading further reduces response time.

1.4.2 4.8 TB cache

DS8A00 supports a maximum 4.8 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 faster response time.

1.4.3 32 Gbps host adapters

The DS8A00 supports 32 Gbps host adapters. This connectivity reduces latency and provides faster single-stream and per-port throughput. These adapters can be configured 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 HBA supports both IBM Fibre Channel Endpoint Security authentication and line-rate encryption. So, the adapter handles this type of data-in-flight encryption autonomously. For IBM Z, the plan is to make it a standard in one of the upcoming zSystems host generations.

1.4.4 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 a maximum of 4.8 TB overall cache sizes, the tables to maintain the cache segments can 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 which data to store in cache that is based on the frequency needs of the hosts

Adaptive Multi-stream Prefetching

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

Intelligent Write Caching (IWC) improves performance through better write-cache management and by 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.

1.4.5 Flash storage

Flash drives sharply improve I/O transaction-based performance over traditional HDDs in standard drive enclosures and deliver a much better latency. Their better IOPS, lower power consumption, less heat generation, and lower acoustical noise offer compelling advantages.

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. Several types and tiers of flash are available. Also, within the higher durability range of enterprise flash, flash modules with different performance characteristics or extra functionality are offered to cover more cost-economical large-capacity use cases and highly-performing and intense workloads with heavy access densities.

For all these scenarios, the DS8A00 offers the different tiers of drives including two performance classes of industry-standard drives and another tier on the IBM FlashCore Modules (FCMs), which features hardware compression and encryption. Each FCM has extra added processing power and DRAM to perform the compression task. Also, caching and internal SLC/QLC tiering is available, which helps to keep latency at the same low levels.

All flash devices utilize the NVM Express (NVMe) protocol, running on the PCIe hardware interfaces. NVMe was developed for SSD flash, and NAND flash is used in our devices. Compared to earlier SAS standards, NVMe allows a much higher I/O parallelism and more transactions and IOPS per drive.

Considering the RAID implementation that is always needed, 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). The transition to PCIe Gen-4 for the HPFE and drive backend has significantly enhanced performance that results in approximately double the typical per-HPFE and per-array throughput compared to the previous DS8000 generation.

The flash drives use the flash RAID adapters in the I/O enclosures, and PCIe connections to the processor complexes. The high-performance flash drives that are available with the Capacity Tier 1 are high-IOPS class enterprise storage devices that are targeted at most I/O-intensive workload applications, which need high-level, fast-access storage. The high-capacity flash drive types of Capacity Tier 2 have a more economical acquisition cost point. For many workloads, the Tier 2 drives can often fulfill standard enterprise workload requirements as a single tier. The Compression Tier gives some extra capacity advantage while maintaining a comparable excellent level of low latency. With its always-on hardware compression due to the large and extra capacities there it gives another price-competitive edge even when compared to Capacity Tier 2.

1.4.6 Performance for IBM Z

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

- With PAVs, a single IBM Z server can 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.

- With HyperPAV, applications can achieve equal or better performance than with PAV alone, use fewer unit control blocks (UCBs), and eliminate the latency in targeting an alias to a base. With HyperPAV, the system can react immediately to changing I/O workloads.
- With SuperPAV, z/OS can use an *alias address* from another logical control unit (LCU).
- Multiple Allegiance expands the simultaneous logical volume access capability across multiple IBM Z servers. With this function, and with PAV, the DS8000 can process more I/Os in parallel, which improves performance and enables greater use of large volumes.
- With I/O priority queuing, the DS8000 series can 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.
- With zHPF, multiple channel commands can 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 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, which improves transactional response times and log throughput. The MM primary volume must 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 bays of the DS8000. Response times in the range of 16 µs are possible for qualifying I/Os. zHyperLink supports read and write I/O.
- With the new zHyperLink Acceleration in DS8A00 and because of the respective PCIe Gen-4 adapters, latency, read, and write throughput have improved. The improvement in write throughput is approximately a factor of 2.5. Further zHyperLink exploitation is possible because of an increased maximum zHL read block size to 32K with a default of 16K, which allows new zHL usage scenarios for customers with larger CI sizes. IMS VSAM databases and larger Db2 pages sizes up to 32K can now use it. And for Writes, MQ (9.4+) log writes can use zHL also, sharply increasing their throughput.
- Writes of zHyperLink 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 µs, that same write operation takes only 27 µs over zHyperLink.
- In an MM replication configuration, Consistent Read from Secondary volumes (CRS) is possible over zHyperLink or zHPF in a HyperSwap configuration.
- With the TCT function, the DS8000 CPCs can be scheduled to directly move data over to tape or to Cloud Object Storage without data going through the host, which can sharply reduce IBM Z MIPS usage during the backup periods and reduce expenses and simplify data archiving operations. The TCT function enables multi-cloud support where up to eight cloud targets can be defined for each DS8000. For more information, see *IBM DS8000 Transparent Cloud Tiering*, SG24-8381.

For more information about performance on IBM Z, see I*BM DS8000 and IBM Z Synergy*, REDP-5186.

1.4.7 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

2

IBM DS8A00 hardware components and architecture

This chapter describes the hardware components of the IBM DS8A00. It provides insights into the architecture and individual components.

This chapter covers the following topics:

- “Flash drive terminology of the DS8A00” on page 40
- “DS8A00 configurations and models” on page 42
- “DS8A00 architecture overview” on page 55
- “I/O enclosures and adapters” on page 61
- “Flash drive enclosures” on page 67
- “Power and cooling” on page 69
- “Hardware Management Console and network” on page 75

© Copyright IBM Corp. 2025. All rights reserved.

2.1 Flash drive terminology of the DS8A00

This section describes the naming conventions that are used to describe the DS8A00 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 DS8A00 (base frame plus optional expansion frame).

Base frame

The DS8A00 has four available base frame models. The model numbers, DS8A80, DS8A50 and DS8A10, depend on the hardware configuration for each. In this chapter, the DS8A00 family name, or model number, are used interchangeably. Table 2-1 lists each of the frame models.

Table 2-1 DS8A00 frame models and expansion frames

DS8A00	Base frame model	Expansion frame model	Max expansion frames
DS8A80	A08	E05	1
DS8A50	A05	E05	1
DS8A10	A01	N/A	N/A
DS8A10	A00	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 DS8A80 model A08 or any DS8A50 model A05.

For more information about the base frame configuration, see 2.2.5, “DS8A00 base frames” on page 49.

For more information about the DS8A10 model A00 Rack-Mounted system, see 2.2.4, “DS8A10 Rack-Mounted configuration” on page 47.

Expansion frame

The DS8A80 and DS8A50 support one optional expansion frame, which provides space for extra storage capacity, and also supports up to two additional I/O enclosure pairs.

With this model, you can place the expansion frame a maximum of 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, “DS8A00 expansion frame” on page 52.

All DS8A00 system memory and processor upgrades can be performed concurrently.

2.1.2 Hardware Management Console

Each DS8A00 contains two IBM Power Systems Hardware Management Consoles (HMC) type 7063-CR2 and two rack control hubs (Ethernet switches) for internal network communications to the central processor complexes (CPCs). Each DS8A00 also includes I/O enclosures and intelligent Power Distribution Units (iPDUs). Clients connect Ethernet cables to both HMCs for remote management access to the DS8A00. A monitor and keyboard with trackpad is installed in the base frame for local management.

Figure 2-1 on page 41 shows the front view of both Power9 HMCs.

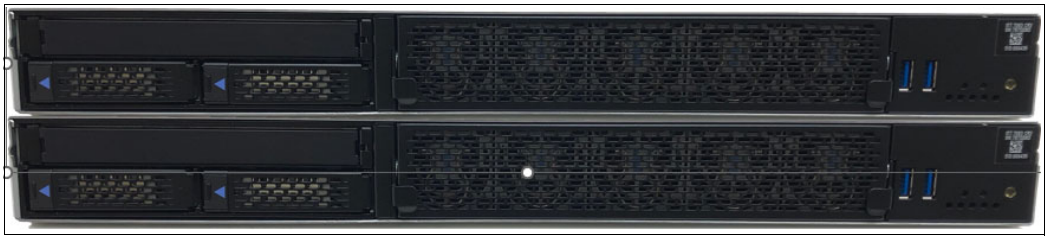


Figure 2-1 Power9 Hardware Management Console (HMC)

2.1.3 Central processor complex

The DS8A00 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 55.

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

Table 2-2 DS8A00 processor and System Memory details

Model	Processor per storage system	System Memory per storage system
A08	Forty-eight cores	4864 GB
A05	Forty cores	1024 GB
		2048 GB
		3584 GB
A01	Twenty cores	256 GB
		512 GB

Both CPCs in a DS8A00 system share the system workload. The CPCs are redundant, and either CPC can fail over to the other for scheduled maintenance, for 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 DS8A00 configurations and models

This section presents the DS8A00 configurations and models at the time of writing. The DS8A00 consists of one base frame, and for the DS8A50, 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.

The DS8A00 storage systems use machine type 5341. Warranty and services are offered as part of Expert Care. Options range from a 1-year base warranty to a 5-year Expert Care Advanced or Premium, where Premium is the default option in eConfig.

2.2.1 DS8A80 configuration

The DS8A80 model 5341-A08 includes two CPCs, each featuring 24-core processors. The base frame supports up to 16 host adapters for both Fibre Channel (FC) and FICON connectivity. Achieving the maximum configuration requires an E05 expansion frame.

Figure 2-2 shows minimum and maximum DS8A80 configuration:

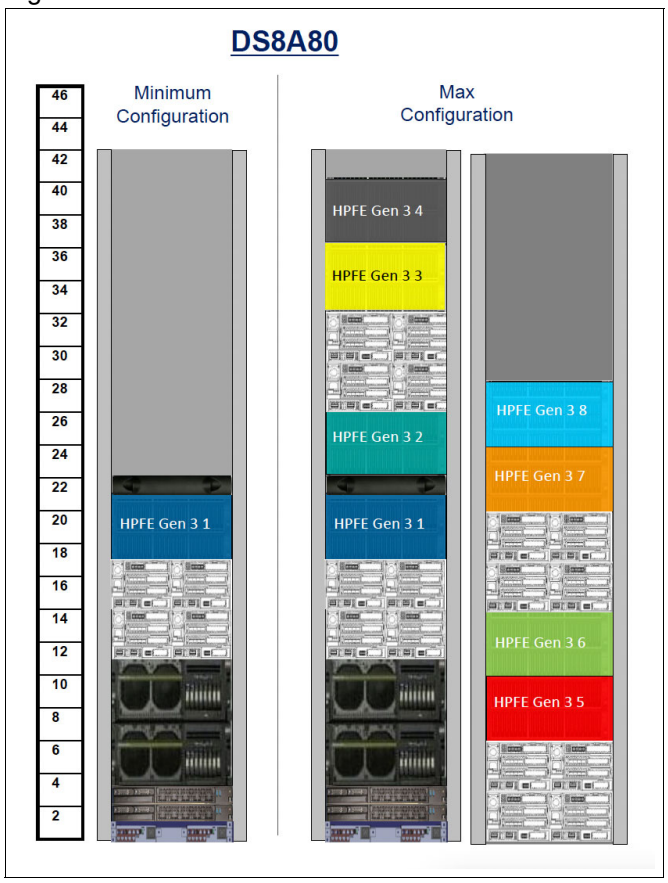


Figure 2-2 Minimum and maximum DS8A80 configuration

Table 2-3 lists the hardware and the minimum and maximum configuration options for the DS8A80 model A08.

Table 2-3 DS8A80 model A08 component list

Features	DS8A80
Rack size	19” 600x1050, 42U fixed
Frames min / max	1 / 2
CPC	Two IBM Power9+ MTM 9009-42G systems
I/O Enclosure pairs min / max	1 / 4
Power9 cores per system min / max	24 / 48
System memory min / max	4864 GB
System NVS memory min / max	256 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 Gen3 pairs min / max	1 / 8
Device adapter pairs min / max	1 / 8
iPDU min / max	2 / 8
Power	Single-phase or three-phase
HMC	Two Power 9 systems (7063-CR2)
Ethernet switches	Two 8-port switches in Base Frame, two additional 24-port switches with an E05 expansion frame
Keyboard/Monitor	1

a. For more information, see 2.4.2, “I/O enclosure adapters” on page 63.
b. For more information, see *Getting Started with IBM zHyperLink for z/OS*, REDP-5493.

Note: The DS8A00 hardware uses iPDUs, non-volatile dual inline memory modules (NVDIMMs) and Backup Power Modules (BPMs).

2.2.2 DS8A50 configuration

The DS8A50 model 5341-A05 is equipped with two CPCs, each with 20-core processors, with a maximum of 16 host adapters in the base frame that support both Fibre Channel (FC) and FICON and a maximum of 16 FC / FICON host adapters in the expansion frame. Maximum configuration requires E05 expansion frame.

Figure 2-3 shows minimum and maximum DS8A50 configuration:

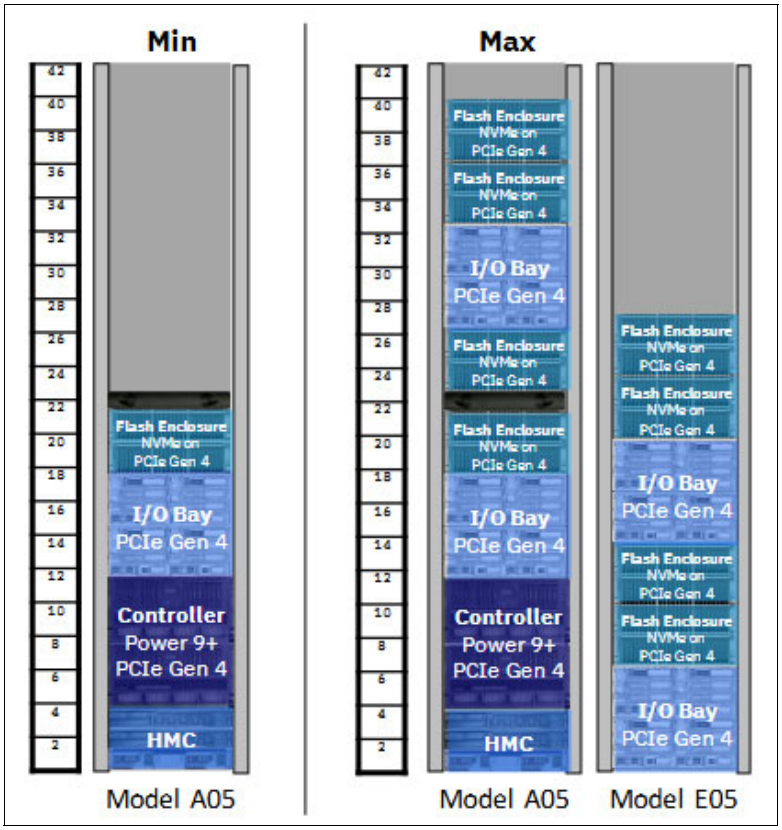


Figure 2-3 Minimum and maximum DS8A50 configuration

Table 2-4 lists the hardware and the minimum and maximum configuration options for the DS8A50 model A05.

Table 2-4 DS8A50 model A05 component list

Features	DS8A50
Rack size	19” 600x1050, 42U fixed
Frames min / max	1 / 2
CPC	Two IBM Power9+ MTM 9009-42G systems
I/O Enclosure pairs min / max	1 / 4
Power9 cores per system min / max	20 / 40
System memory min / max	1024 GB / 3584 GB
System NVS memory min / max	64 GB / 192 GB
Host adapters min / max[a]	2 / 32
Host adapter ports min / max	8 / 128
zHyperLink adapters min / max[b]	0 / 10
Flash drives min / max	16 / 384

Features	DS8A50
HPFE Gen3 pairs min / max	1 / 8
Device adapter pairs min / max	1 / 8
iPDU min / max	2 / 8
Power	Single-phase or three-phase
HMC	Two Power 9 systems (7063-CR2)
Ethernet switches	Two 8-port switches in Base Frame, two additional 24-port switches with an E05 expansion frame
Keyboard/Monitor	1

a. For more information, see 2.4.2, “I/O enclosure adapters” on page 63.

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

2.2.3 DS8A10 Rack configuration

The DS8A10 model 5341-A01 is equipped with two CPCs, each with 10-core processors, with a maximum of 16 FC / FICON host adapters and a maximum of four zHyperLink adapters.

Figure 2-4 shows minimum and maximum configuration of a DS8A10-A01 system.

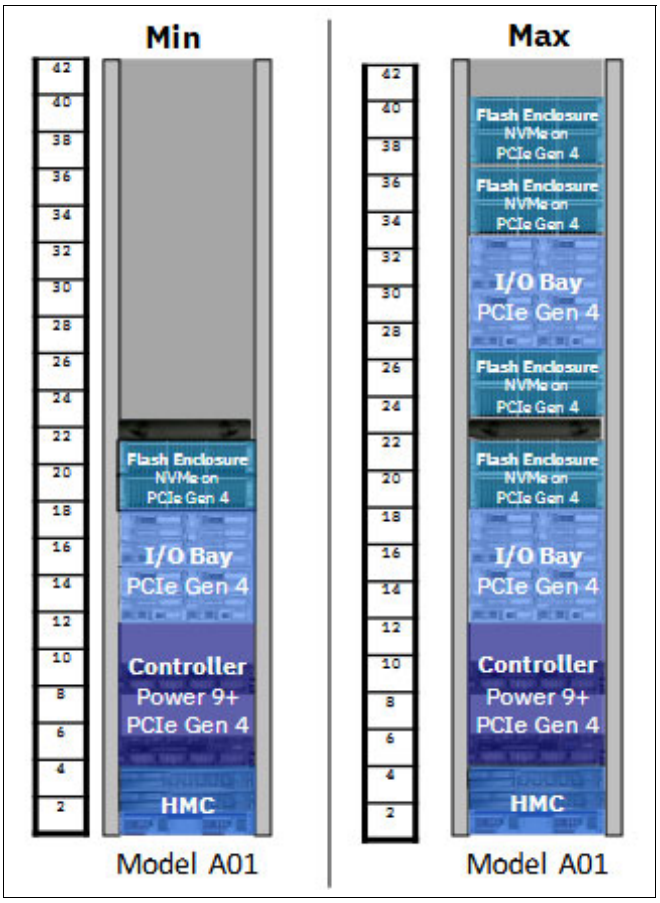


Figure 2-4 DS8A10 model A01 system min. and max. configuration

Table 2-5 lists the hardware and the minimum and maximum configuration options for the DS8A10 model A01.

Table 2-5 DS8A10 model A01 component list

Features	DS8A10
Rack size	19” 600x1050, 42U fixed
Frames	1
CPC	Two IBM Power MTM 9009-22G systems
I/O bay pairs min / max	1 / 2
Power9 cores per storage system	20
System memory min / max	256 GB / 512 GB
System non-volatile storage (NVS) memory min / max	16 GB / 32 GB

Features	DS8A10
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 Gen3 pairs min / max	1 / 4
Device adapter pairs min / max	1 / 4
iPDU min / max	2 / 4
Power	Single-phase or three-phase
HMC	Two Power9 systems (7063-CR2)
Ethernet switches	Two 8-port switches
Keyboard/Monitor	1

a. For more information, see 2.4.2, “I/O enclosure adapters” on page 63.
b. For more information, see *Getting Started with IBM zHyperLink for z/OS*, REDP-5493.

2.2.4 DS8A10 Rack-Mounted configuration

The DS8A10 Rack-Mounted model A00 provides a modular rack-mountable enterprise storage system within DS8A00 family.

The DS8A10 model A00 can be integrated into existing client-provided racks which conform to EIA-310-D specification. The model A00 uses the same hardware components that are found in the rack-based DS8A00 systems and offers all the advanced features while reducing rack space, data center footprint and cost. The DS8A10 model A00 can be ordered with a DS8A00 rack if required; if ordered with this rack, placement of components in the rack is done at the data center.

The DS8A10 model A00 is equipped with two CPCs, each with 10-core processors, with a maximum of 16 FC / FICON host adapters and a maximum of four zHyperLink adapters.

Figure 2-5 shows minimum and maximum configuration of a DS8A10-A00 system.

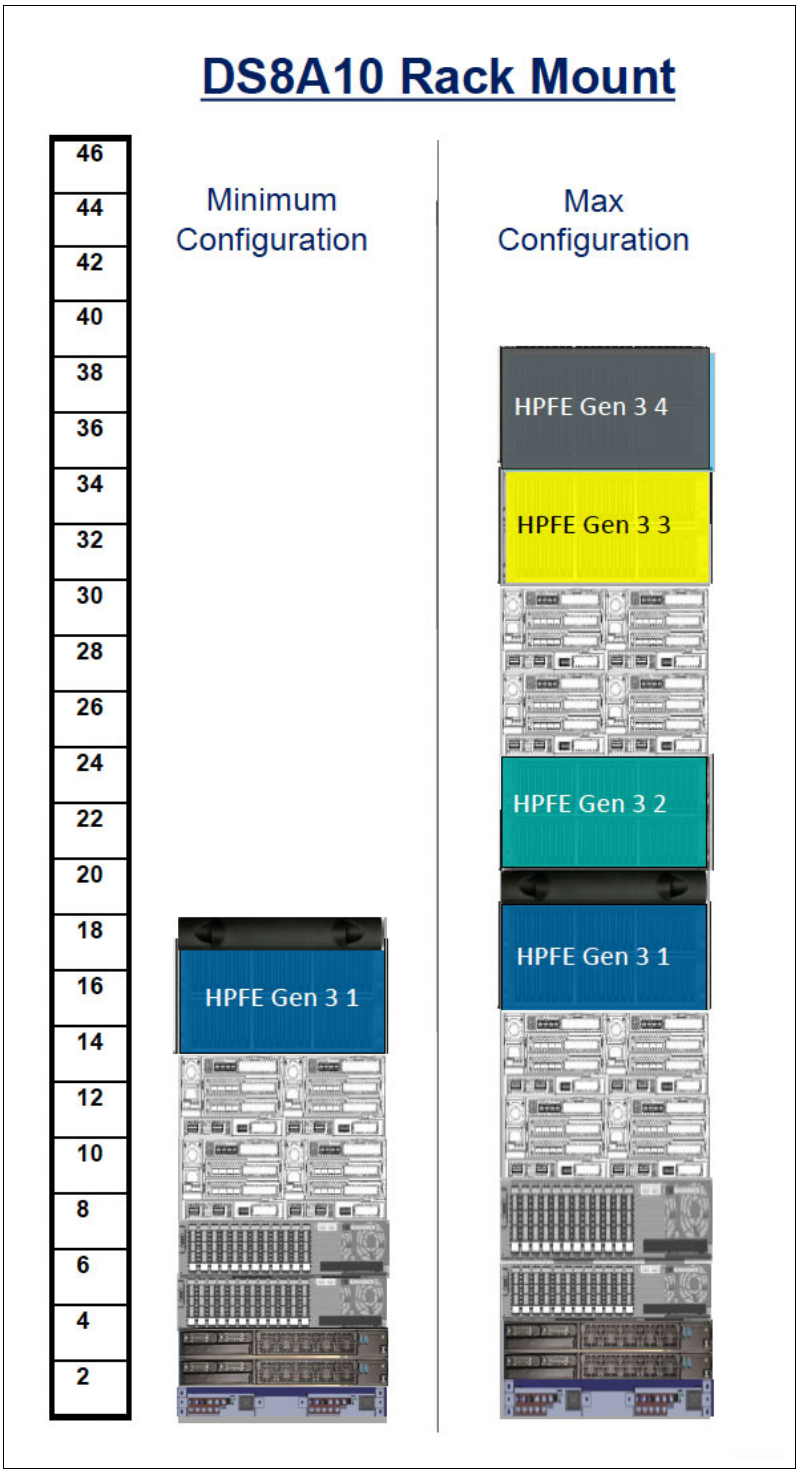


Figure 2-5 DS8A10 model A00 system min. and max. configuration

Table 2-6 on page 49 lists the hardware and the minimum and maximum configuration options for the DS8A10 model A00.

Table 2-6 DS8A10 model A00 component list

Features	Features
CPC	Two IBM Power 9+ MTM 9009-22G systems
I/O bay pairs min / max	1 / 2
Power9 cores per storage system	20
System memory min / max	256 GB / 512 GB
System non-volatile storage (NVS) memory min / max	16 GB / 32 GB
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 Gen3 pairs min / max	1 / 4
Device adapter pairs min / max	1 / 4
Customer provided or iPDU min / max	2 / 4
Power	Single-phase or three-phase
HMC	Two Power9 systems (7063-CR2)
Ethernet switches	Two 8-port switches
Keyboard/Monitor	1

a. For more information, see 2.4.2, “I/O enclosure adapters” on page 63.
b. For more information, see *Getting Started with IBM zHyperLink for z/OS*, REDP-5493.

2.2.5 DS8A00 base frames

As mentioned in 2.1.1, “Storage system” on page 40, the DS8A00 is available in different rack 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 40.

The DS8A00 base racks accommodate the following components:

- Up to four HPFE Gen3 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 Gen3 enclosure pair can accommodate 16, 32, or 48 flash drives.

  The DS8A00 supports Industry Standard NVMe drives and IBM FlashCore Modules (FCM). The Industry Standard NVMe drives are available in 1.6 TB, 1.92 TB, 3.84 TB, and 7.68 TB capacities. The IBM FlashCore Models which support compression are available in 4.8 TB, 9.6 TB, and 19.2 TB capacities. For more information about FCMs, see IBM FlashCore Module Product Guide.

  All HPFE Gen 3 enclosures have redundant power and integrated cooling, which draws air from front to rear. For more information about cooling requirements, see Chapter 5, “IBM DS8000 physical planning and installation” on page 143.

- The rack power subsystem

  The DS8A00 family 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 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 Gen3 pair is 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 69.

- Hardware Management Consoles

  Each base frame contains two Power9 HMCs in all DS8A00 models. Clients connect Ethernet cables to both HMCs for remote management access to the DS8A00. A For more information about the HMC, see 2.7, “Hardware Management Console and network” on page 75.

  Each base frame includes a 1U keyboard or display tray that is used for local management through both HMCs.

- Power9 processor-based CPCs

  Each base frame accommodates two Power9+ servers, which support PCIe Gen4, also known as CPCs. The Power9+ servers contain the processors and memory that drive all functions to operate the DS8A00 storage facility image (SFI). System memory and processor cores in the DS8A00 can be upgraded concurrently. The DS8A50 configurations run on two 4U IBM Power System S924 servers. The DS8A10 systems use two 2U IBM Power Systems S922 servers. For more information about the CPCs, see 2.3.1, “IBM Power9 processor-based CPCs” on page 55.

  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 58.

- I/O enclosures

  The DS8A00 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 Gen4 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 32 Gbps encryption-capable PCIe Gen3 FC / FICON host adapters in each I/O enclosure, supporting either SW or LW.

    Note: Intermix of SW and LW adapters is also allowed. For HA, IBM recommends installing host adapters in pairs. The 32 Gbps HA adapter is using the same technology as in DS8900F systems but with a different form factor.

    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.
    - FCP and FICON are not supported simultaneously on the same port.

Note: The Fibre Channel Arbitrated Loop (FC-AL) topology is not supported on DS8A00 host adapters.

- Device adapters (DAs):
  - One Device adapter pair is required for each HPFE Gen3 pair.
  - Each I/O enclosure pair supports up to two DA pairs.
  - DA pairs are connected to HPFE Gen3 pairs through redundant NVMe over PCIe Gen4 connections.
- zHyperLink connections to IBM Z hosts:
  - Supports direct connectivity to IBM Z at distances up to 150 m.
  - Install the zHyperLink adapters in pairs in two different I/O enclosures.
  - zHyperLink cables are available in lengths of 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 DS8A00 models have zHyperLink capability.

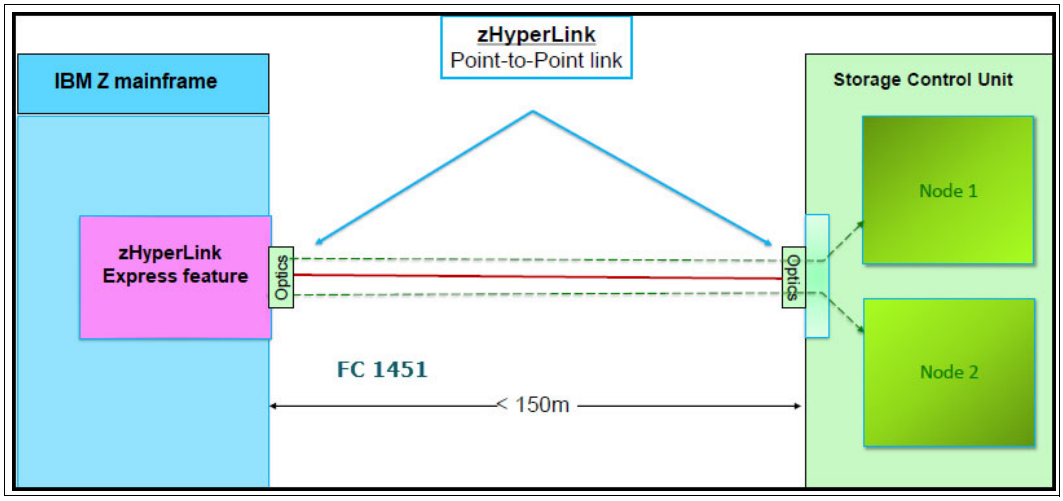


Figure 2-6 DS8A00 zHyperLink connection to the system

- 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 2-18 on page 64.
- The number of zHyperLink adapters that can be installed into the DS8A00 system depends on the number of cores per CPC. The supported combinations for zHyperLink port availability are shown in Table 2-7 on page 52.

Table 2-7 zHyperLink availability for DS8A00 models

System or model	Cores per CPC (DS8A00 server)	zHyperLink support	Max zHyperLink connections (increments of 2)
DS8A80 Base Frame model A08	24	Yes	8
DS8A80 model Expansion Frame E05		Yes	12
DS8A50 Base Frame model A05	20	Yes	8
DS8A50 model Expansion Frame E05		Yes	8
DS8A10 model A01	10	Yes	4
DS8A10 model A00	10	Yes	4

Note: A maximum of 12 zHyperLink connections are supported in a DS8A80 system.

For more information about I/O enclosures and I/O adapters, see 2.4, “I/O enclosures and adapters” on page 61.

2.2.6 DS8A00 expansion frame

Both the DS8A80 and the DS8A50 configuration support an optional model E05 42U high-end expansion frame.

The DS8A80 and DS8A50 systems support the installation of an expansion frame without any additional features.

Expansion frame components

The model E05 expansion frame accommodates the following components:

- Up to four HPFE Gen3 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 enclosure pair can accommodate 16, 32, or 48 flash drives.

  The DS8A00 supports Industry Standard NVMe drives and IBM FlashCore Modules (FCM). The Industry Standard NVMe drives are available in 1.6 TB, 1.92 TB, 3.84 TB, and 7.68 TB capacities. The IBM FlashCore Models, which support hardware compression, are available in 4.8 TB, 9.6 TB, and 19.2 TB capacities. For more information about FCMs, see IBM FlashCore Module Product Guide.

  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 DS8000 physical planning and installation” on page 143.

- The rack power subsystem

  The DS8A00 family 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 organized and installed in pairs. One iPDU pair is installed by default in each frame. For more information about the iPDUs, see 2.6, “Power and cooling” on page 69.

- I/O enclosures

  Both the DS8A80 and DS8A50 model E05 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 Gen4 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 32 Gbps FC / FICON encryption-capable PCIe Gen3 host adapters in each I/O enclosure, supporting either SW or LW.

    Note: Intermix of SW and LW adapters is also allowed. For HA, IBM recommends installing host adapters in pairs. The 32 Gbps HA adapter uses the same technology that is used in DS8900F systems but with a different form factor.

    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
    - FCP and FICON are not supported simultaneously on the same port.

  – Device adapter:

    - One device adapter pair is required for each HPFE Gen3 pair.
    - Each I/O enclosure pair supports up to two DA pairs.
    - DA pairs are connected to HPFE Gen3 pairs through redundant NVMe over PCIe Gen4 connections.

  – zHyperLink connections to IBM Z hosts:

    - Supports direct connectivity to IBM Z at distances up to 150 m.
    - zHyperLink adapters should be installed in pairs in two different I/O enclosures.

For more information, see 2.4, “I/O enclosures and adapters” on page 61.

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 54, 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.

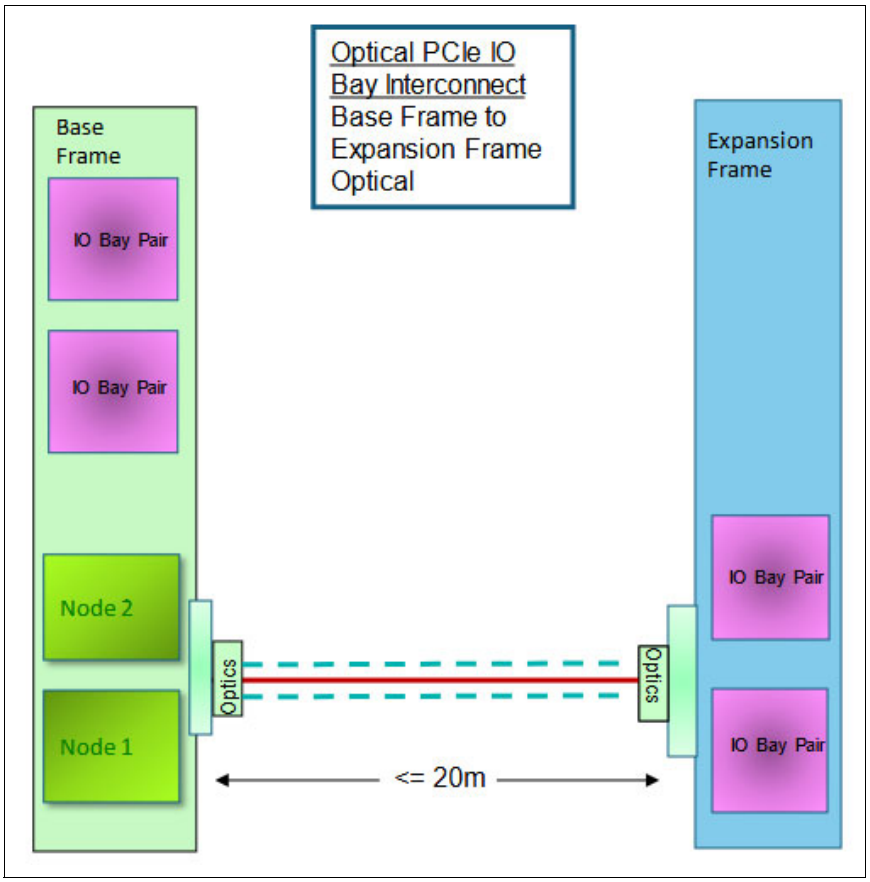


Figure 2-7 Expansion frame at a distance

2.2.7 Scalable upgrades

The hardware features that are supported in the DS8A00 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 2-10 on page 57. Each of the DS8A00 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 DS8A00 functions require a license key. For more information about licensed functions, see Chapter 7, “IBM DS8000 G10 features and licensed functions” on page 195.

2.3 DS8A00 architecture overview

This section provides an architectural overview of the major components of the DS8A00:

- 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 DS8A00 systems include two CPCs:

- In the DS8A80 and DS8A50 configurations, the CPCs are IBM Power 9009-42G servers, which have two processor sockets. The processors run in Turbo mode and typically achieve cycle speeds of 3.8 GHz.

  For the DS8A80 configuration, the CPCs are populated with two 12-core processors, for a combined total of 24 cores per CPC. The DS8A80 systems support a maximum of 2432 GB of system memory per CPC

  For the DS8A50 configuration, the CPCs are populated with two 10-core processors, for a combined total of 20 cores per CPC. The DS8A50 systems support a maximum of 1792 GB of system memory per CPC.
- The CPCs that are used in both the DS8A80 and DS8A50 feature the following configuration:
  - Two single-chip module (SCM) sockets
  - Thirty-two DDR4 dual inline memory module (DIMM) slots
  - Two BPMs
  - Five x16 Gen 4 full-height half-length (FHHL) slots
  - Two x8 Gen 4 FHHL slots with x16 connectors
  - Four x8 Gen 4 FHHL slots
  - Two 800 GB NVMe U.2 drives
  - Four power supply units (PSUs).

Figure 2-8 shows a CPC front view for the DS8A50 systems.



Figure 2-8 DS8A50 CPC: front view

- In both the DS8A10 model A01 and model A00 configurations, the CPCs are IBM Power 9009-22G servers, which are populated with two 8-core processors with 5 cores used per processor, which makes 10 cores in total per CPC. The processors run in Turbo mode and typically achieve cycle speeds of 3.8 GHz.

  The DS8A10 supports a maximum of 256 GB of system memory per CPC.
- The CPCs that are used in DS89A0 systems feature the following configuration:
  - Two SCM sockets
  - Thirty-two DDR4 DIMM slots
  - One BPM
  - Five x16 Gen4 low-profile, half-length slots
  - Two x8 Gen4 low-profile, half-length slot (with x16 connectors)
  - Four x8 Gen4 low-profile, half-length slots
  - Two 800 GB NVMe U.2 drives
  - Two PSUs.

Figure 2-9 shows the CPC front view as configured in the DS8A10 system.

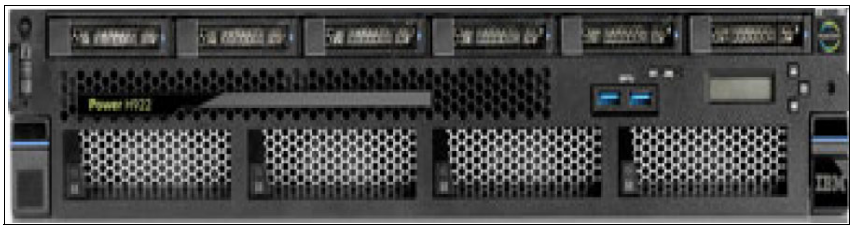


Figure 2-9 DS8A10 CPC: front view

For more information about the server hardware that is used in the DS8A00, see *IBM Power Systems S922, S914, and S924 Technical Overview and Introduction Featuring PCIe Gen 4 Technology*, REDP-5595.

In the DS8A00, 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 2-10 on page 57 shows the supported components for the DS8A00 processor and memory options. NVS values are typically 1/16th of installed system memory, except for the smallest systems with 256 GB system memory, where only 16 GB, that is, 8 GB per CPC, is used as NVS, and for the biggest systems of 4.8 TB memory, where NVS remains at 256 GB.

Model	Processors per System	System Memory (GB)	System NVS (GB)	Expansion Frame	Min / Max Host Adapters	Min / Max zHyperLink	Min / Max Flash Enclosures	Min / Max Flash Drives
A08 Multiframe	48-cores	4864	256	0 to 1	2 (8 ports) / 32 (128 ports)	0 / 12	1 / 8	16 / 384
A05 Multiframe	40-cores	1024 2048 3584	64 128 192	0 to 1	2 (8 ports) / 32 (128 ports)	0 / 10	1 / 8	16 / 384
A01 Single Frame	20-cores	256 512	16 32	n/a	2 (8 ports) / 16 (64 ports)	0 / 4	1 / 4	16 / 192
A00 Rack Mount	20-cores	256 512	16 32	n/a	2 (8 ports) / 16 (64 ports)	0 / 4	1 / 4	16 / 192

Figure 2-10 Supported components for the DSA800 processor and memory options

2.3.2 Processor memory

The DS8A80 configuration offers 4864 GB of total system memory. DS8A50 configuration offers up to 3584 GB of total system memory. The DS8A10 Racked and Rack-Mountable 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 DS8A00 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 DS8A00 configuration options are based on the total installed memory, which in turn depends on the number of installed and active processor cores:

- The DS8A80 is preinstalled with 2432 GB per server.
- In the DS8A50, you can install a memory upgrade from 512 GB per server to 1024 GB per server or to 1792 GB per server nondisruptively.
- In the DS8A10 models A00 and A01, you can install a memory upgrade from 128 GB per server to 256 GB per server.

Caching is a fundamental technique for reducing I/O latency. Like other modern caches, the DS8A00 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.

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 2-11 on page 59. The bottom two RJ45 copper ports can be used for Transparent Cloud Tiering (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 DS8A00 systems. The optional 10 Gbps Ethernet adapter is in the P1-C10 slot in DS8A00 systems. For more information about TCT, see *IBM DS8000 Transparent Cloud Tiering*, SG24-8381.

Figure 2-11 shows the location codes of the CPCs in DS8A80 and DS8A50 systems. Figure 2-12 shows the location codes of the CPC in a DS8A10 systems.

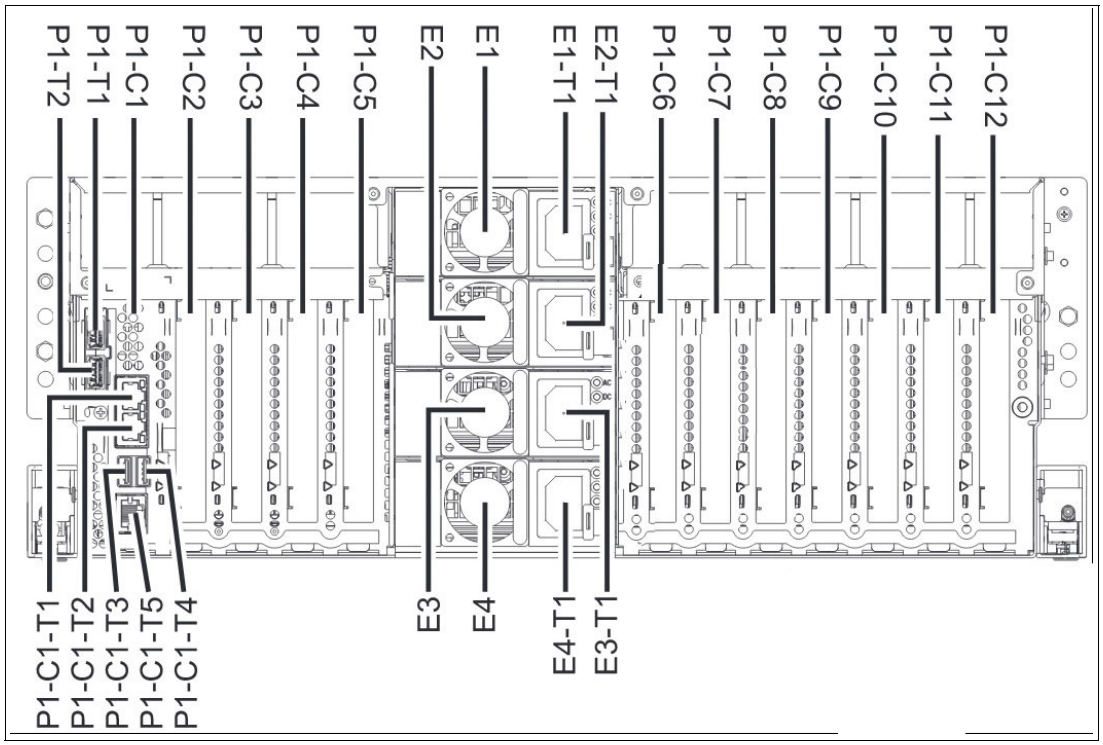


Figure 2-11 Location codes of the CPC in DS8A80 and DS8A50 systems in the rear

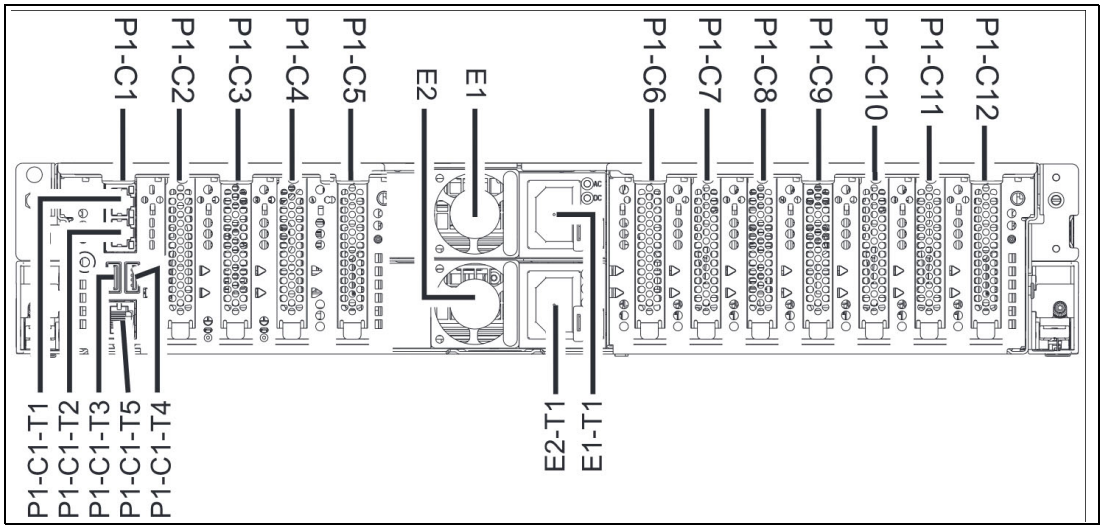


Figure 2-12 Location codes of the CPC in the DS8A10 systems in the rear

2.3.5 Peripheral Component Interconnect Express adapters

Each DS8A00 CPC contains multiple PCIe Gen4 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 DS8A00 CPC.

A DS8A00 CPC is equipped with the following PCIe Gen 4 adapters:

- The DS8A80 and DS8A50 systems have 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 DS8A10 systems have four single-port PCIe adapters in slots C6, C8, C9, and C12.

Figure 2-13 shows the PCIe adapter locations in the DS8A50 CPC. Figure 2-14 shows the PCIe adapter locations in the DS8A10 CPC.

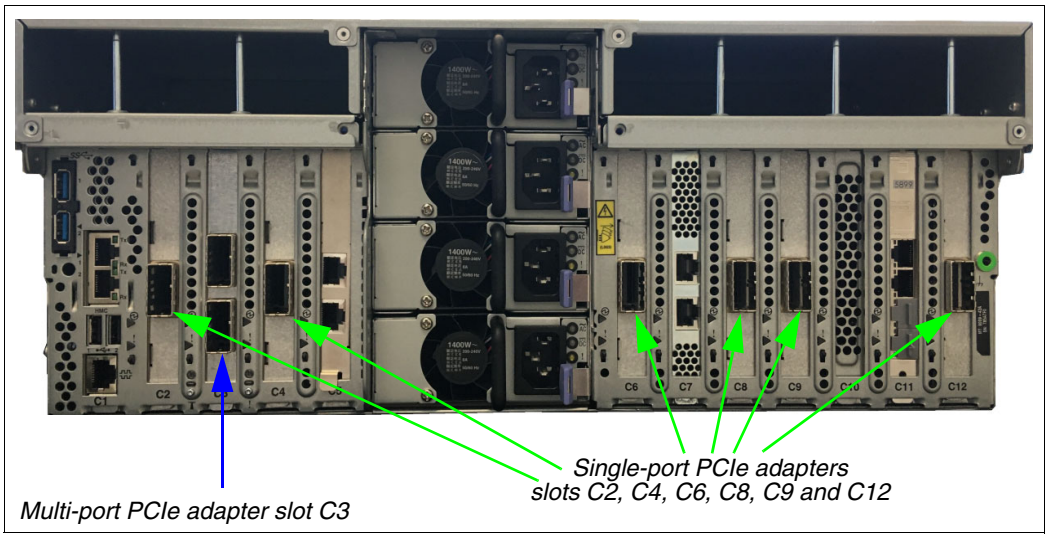


Figure 2-13 PCIe adapter locations in the DS8A80 and DS8A50 CPC: rear view

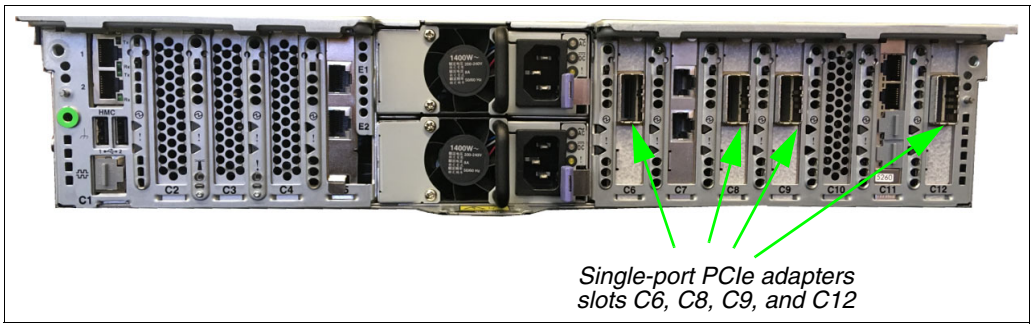


Figure 2-14 PCIe adapter locations in the DS8A10 CPC: rear view

2.4 I/O enclosures and adapters

The DS8A00 base frame and expansion frame (if installed) contain I/O enclosures, which are installed in pairs.

The I/O enclosures are PCIe Gen4-capable, and are attached to the CPCs with 8-lane PCIe Gen4 cables. The I/O enclosures have six PCIe adapter slots, plus two CXP connectors.

- Both the DS8A80 and DS8A50 CPCs have up to six 1-port and one 2-port PCIe Gen4 adapters that provide connectivity to the I/O enclosures.
- DS8A10 CPCs have up to four 1-port PCIe Gen4 adapters that provide connectivity.

Figure 2-15 on page 62 shows the DS8A10 CPC-to-I/O enclosure connectivity.

The DS8A50 configuration requires no extra features to support an expansion frame.

One or two I/O enclosure pairs can be installed in the base frame of the DS8A00 and also in the E05 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 up to two zHyperLink adapters. For more information about zHyperLink availability for DS8A00 models, see Table 2-7 on page 52.

Each I/O enclosure has the following characteristics:

- 3U rack-mountable enclosure
- Eight PCIe Gen4 slots
  - four for host adapters
  - two for device adapters
  - two zHyperLink adapters
- Two PCIe connections to the CPCs
- Redundant power and cooling

Figure 2-15 on page 62 shows the DS8A80 and DS8A50 CPC-to-I/O enclosure connectivity.

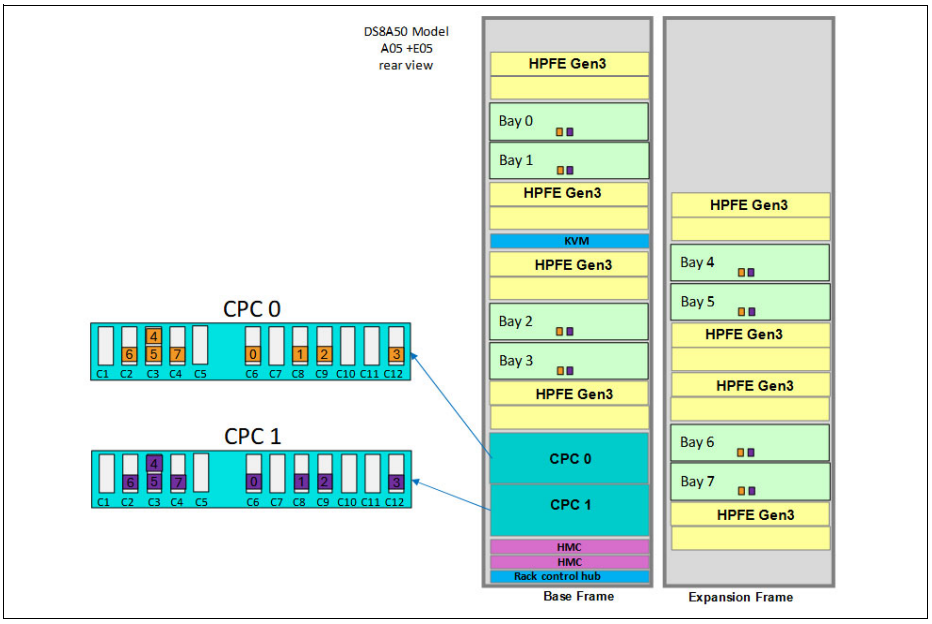


Figure 2-15 DS8A80 and DS8A50 I/O enclosure connections to the CPCs

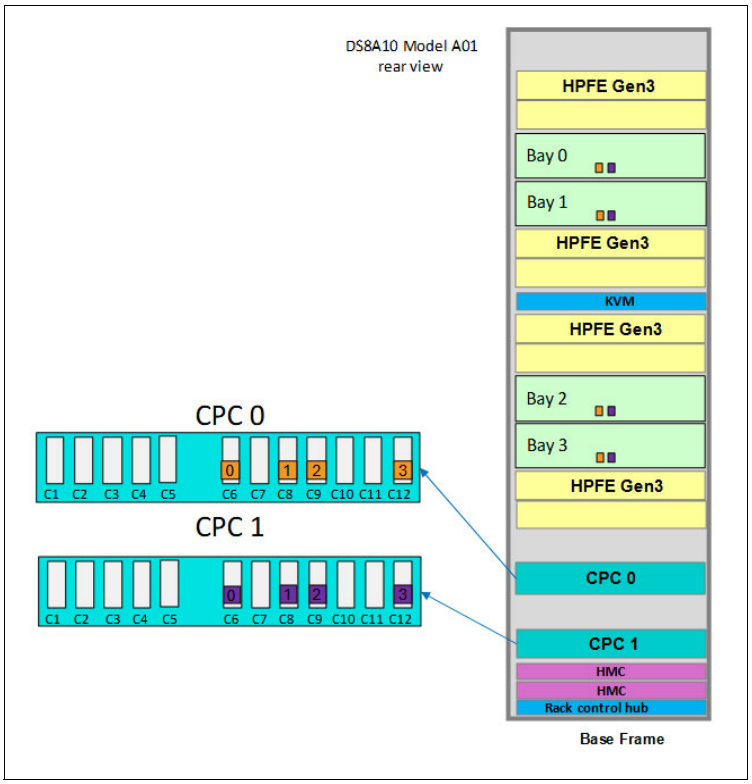


Figure 2-16 DS8A10 I/O enclosure connections to the CPC

2.4.1 Cross-cluster communication

The DS8A00 I/O enclosure hardware is shared between the servers (CPC). A high-speed communication between the CPCs is established by using PCIe Gen4 paths through the I/O enclosures. 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.

Figure 2-17 shows PCIe connections of a DS8A80, One CPC is on the left side and one CPC is on the right side.

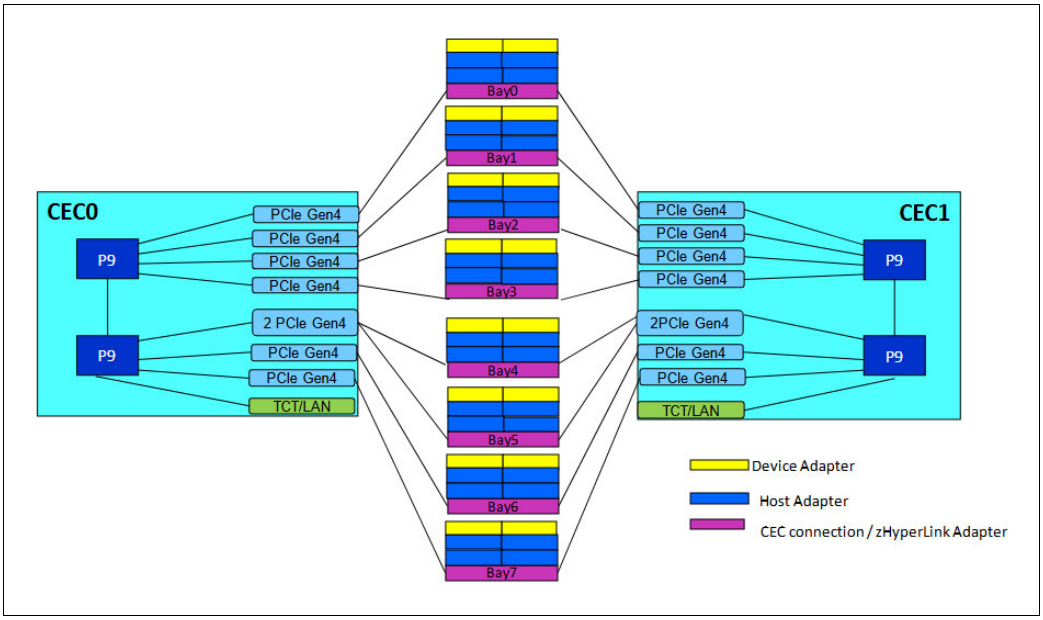


Figure 2-17 DS8A80 series PCIe communications paths

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 DS8A00 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 and two zHyperlink adapters (zHL) that provide attachment to host systems Also, the I/O bay can contain up to two device adapters to provide attachment to the High-Performance Flash Enclosure (HPFE) Gen3. Each I/O bay has two PCIe Gen4 x8 CXP connectors on the I/O bay for the internal PCIe fabric connections to CPC 0 and CPC 1.

Figure 2-18 shows the DS8A00 I/O bay adapter layout.

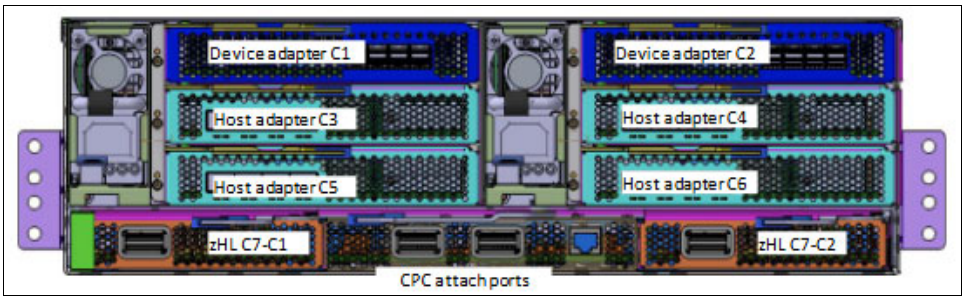


Figure 2-18 DS8A00 I/O bay adapter layout

DS8A00 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.

Both 32 Gbps Short Wave and 32 Gbps Long Wave host adapters are available. Both have four ports and the adapters can auto-negotiate their data transfer rate down to 8 Gbps full-duplex data transfer. Both host adapters support FCP or FICON protocol.

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).

The DS8A80 and DS8A50 configurations support a maximum of 16 host adapters in the base frame and 16 extra host adapters in the model E05 expansion frame. Both DS8A10 models support a maximum of 16 host adapters.

Host adapters are installed in slots 3, 4, 5, and 6 of the I/O enclosure. Figure 2-18 shows the locations for the host adapters in the DS8A00 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 DS8A00 port configurations

Model	Min/Max host adapters	Min/Max host adapter ports	Max zHyperLink adapters
A08/A05	2/16	8/64	8
A08 + E05	2/32	8/128	12
A05 + E05	2/32	8/128	10[a]
A01	2/16	8/64	4

a. Maximum of eight in either the base frame or the expansion frame. Maximum of 10 per system.

Table 2-9 shows the host adapter installation order for the DS8A00 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 expansion frame.

Table 2-9 DS8A00 host adapter installation order Fibre Channel

I/O bay	Slot number					
	C1	C2	C3	C4	C5	C6
For two I/O enclosures (all models, including *Model E05*[a])						
bay 02 / *06*	N/A	N/A	1	3	5	7
bay 03 / *07*	N/A	N/A	2	4	6	8
For four I/O enclosures (Model A05, Model A01, *Model E05*[a])						
bay 00 / *04*	N/A	N/A	3	7	11	15
bay 01 / *05*	N/A	N/A	4	8	12	16
bay 02 / *06*	N/A	N/A	1	5	9	13
bay 03 / *07*	N/A	N/A	2	6	10	14

a. For the DS8A80 and DS8A50 model E05, the enclosure numbers appear in *emphasized* text, and the plug order is the same as the other models.

The DS8A00 uses the 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 DS8A00 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 maximum distances of 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 the distance limitations of a 32 Gbps HA with SW using OM3, OM4,or OM5, see Table 2-10.

Table 2-10 SW link distance

Speed	OM3 link distance	OM4 link distance	OM5 link distance
32 Gbps	70 m (229.6 ft.)	100 m (328 ft.)	150 m (492.13 ft.)

zHyperlink adapters

zHyperlink adapters can be attached via a point-to-point connection to IBM Z servers running the z/OS operating system. It provides a low latency interface to a z/OS LPAR, which allows synchronous I/Os and means the CPU is waiting for a response from the DS8A00. For more

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

In each I/O enclosure two zHyperLink (zHL) adapters can be installed. The number of supported zHL adapters per system depends on the number of processor cores. A maximum of 10 zHL adapters can be installed in a DS850F and 4 zHL adapters are supported in a DS8A10 system. The preferred installation order of the zHL adapters is shown in Table 2-11.

Table 2-11 zHyperLink adapter plug order

zHyperLink Plug order				**Location**	
DS8A00 2 Frame 4x I/O enclosure pairs	**DS8A00 2 Frame 3x I/O enclosure pairs**	**DS8A00 1 Frame 2x I/O enclosure pairs**	**DS8A00 1 Frame 1x I/O enclosure pairs**	**Base Frame location**	**Expansion Frame location**
1	1				Bay 6 C7-C1
2	2				Bay 7 C7-C1
3					Bay 4 C7-C1
4					Bay 5 C7-C1
5	3	1	1	Bay 2 C7-C1	
6	4	2		Bay 1 C7-C1	
7	5	3	2	Bay 3 C7-C1	
8	6	4		Bay 0 C7-C1	
9	7				Bay 6 C7-C2
10	8				Bay 7 C7-C2
11					Bay 4 C7-C2
12					Bay 5 C7-C2
		5	3	Bay 2 C7-C2	
		6		Bay 1 C7-C2	
		7	4	Bay 3 C7-C2	
		8		Bay 0 C7-C2	

Device adapters

Device adapters (DA) provide redundant NVMe access to the internal storage devices and support RAID6 and RAID10. Each DA manages a pair of HPFE Gen3 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 Gen4 network. The DAs are responsible for managing and monitoring the RAID arrays (RAID6 / RAID10) in the HPFE enclosures. The DAs provide remarkable performance because of a high-function and high-performance NVMe design. To ensure maximum data integrity, the adapter supports metadata creation and checking.

2.5 Flash drive enclosures

The DS8A00 is equipped with High-Performance Flash Enclosure (HPFE) Gen3 storage enclosures:

- Flash enclosure pairs connect to devices adapters (DA) by using NVMe connections.
- DAs are installed in the C1 and C2 slots of I/O enclosures.

HPFE Gen3 enclosures are always installed in pairs. Each enclosure pair supports 16, 32, or 48 flash drives. In Figure 2-19, a pair of HPFE Gen3 enclosures is shown.

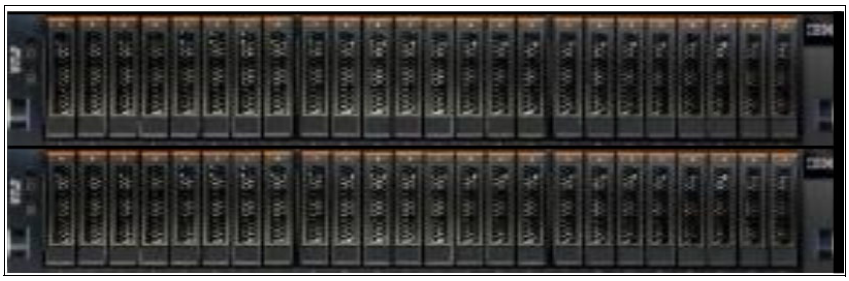


Figure 2-19 HPFE Gen3 enclosure

Each HPFE Gen3 pair is connected to a redundant pair of devices adapters. The PCIe Gen 4 DAs are installed in the DS8A00 I/O enclosures.

Both The DS8A80 and DS8A50 configurations can support up to four HPFE Gen3 pairs in the base frame, and up to four HPFE Gen3 pairs in the expansion frame for a total of eight HPFE Gen3 pairs, with a maximum of 384 flash drives. Both DS8A10 models support configurations up to four HPFE Gen3 pairs, with a maximum of 192 flash drives.

Flash drives in the HPFE Gen3

Each HPFE Gen3 pair can contain 16, 32, or 48 flash drives. HPFE Gen3 supports Industry Standard NVMe drives with 1.6 TB, 1.92 TB, 3.84 TB, and 7.68 TB capacity, and IBM FlashCore Modules with 4.8 TB, 9.6 TB, and 19.2 TB capacity. Both drive types are NVMe drives supporting data-at-rest encryption, and the IBM FlashCore Modules also provide data compression capability. All flash drives in an HPFE Gen3 enclosure pair must be of the same type, either Industry-Standard NVMe or IBM FlashCore Module.

Flash drive sets

Flash drives are ordered in drives sets of 16. The HPFE Gen3 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 help ensure consistent airflow through an enclosure. For HPFE Gen3, one filler feature provides a set of 16 fillers.

Industry Standard NVMe flash drives

The DS8A00 system supports 2.5-inch Industry Standard NVMe flash drives. All flash drives are Full Disk Encryption (FDE) capable, Table 2-12 on page 68 shows supported NVMe flash drives. For more information about licensed features, see Chapter 7, “IBM DS8000 G10 features and licensed functions” on page 195.

Table 2-12 Supported high-performance flash drives

Feature Code	Drive capacity	Drive type	RAID support (Default RAID 6)
AAA0	1.6 TB	2.5-in NVMe	6 and 10
AAB0	1.92 TB	2.5-in NVMe	6 and 10
AAB1	3.84 TB	2.5-in NVMe	6 and 10
AAB2	7.68 TB	2.5-in NVMe	6 and 10

IBM FlashCore Modules

The DS8A00 system is the first DS8000 generation that supports 2.5-inch IBM FCMs which are NVMe drives supporting data-at-rest-encryption and hardware compression. The hardware compression on each FCM uses LZ77 + Dynamic Huffman Encoding on 16 KB blocks. Compression and thin provisioning within FCMs cannot be disabled.

For FCMs, Easy Tier auto mode must be enabled to allow extent management.

FCMs supports space reclamation on a track boundary (56 KB CKD / 64 KB Open) instead of an extent boundary.

Note: With DS8A00 Release 10.1, no intermix of Industry-Standard NVMe and IBM FCMs in a High-Performance Flash Enclosure or in a storage pool is supported.

The IBM FlashCore Modules are available in 4.8 TB, 9.6 TB, and 19.2 TB capacities, see Table 2-13. For more information about FCMs, see IBM FlashCore Module Product Guide.

Table 2-13 IBM FlashCore Modules capacity

Feature Code	Drive capacity	Drive type	RAID support (Default RAID 6)
AAAA	4.8 TB	2.5-in. IBM FCM	6, 10
AAAB	9.6 TB	2.5-in. IBM FCM	6, 10
AAAC	19.2 TB	2.5-in. IBM FCM	6, 10

Arrays and spares

Each HPFE Gen3 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 Gen3 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 Gen3 Pair are 6+P+Q®.

Note: For all drive types, RAID 6 is the default in the DS GUI and DS CLI, but RAID 10 is optional.

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 2-14.

Table 2-14 Maximum usable and provisioned capacity based on system cache size

Cache	Maximum usable size with large extents	Maximum provisioned size with large extents	Maximum usable size with small extents	Maximum provisioned size with small extents
Less than or equal to 512 GB	Fixed-Block (FB): 8192 TiB Count Key Data (CKD): 7304 TiB	FB: 8192 TiB CKD: 7304 TiB	FB: 1024 TiB CKD: 1102 TiB	FB: 2048 TiB CKD: 1826 TiB
Greater than 512 GB	FB: 32768 TiB CKD: 29216 TiB	FB: ~32768 TiB CKD: ~29216 TiB	FB: 4096 TiB CKD: 4410 TiB	FB: ~8192 TiB CKD: ~7304 TiB

Table 2-15 shows the maximum number of flash drives and maximum raw storage capacity for the different models.

Table 2-15 Maximum raw storage capacity per model

Model / Processors	System memory (GB)	Maximum number of flash drives	Maximum raw storage capacity[a]
DS8A80 / 48-core	4864	384	7372.8 TB
DS8A50 / 40-core	1 024 2 048 3 584	384	7372.8 TB
DS8A10 / 20-core	256 512	192	3686.4 TB

a. Using 19.2 TB FCM drives without compression.

2.6 Power and cooling

The DS8A00 power and cooling systems are highly redundant. The components are described in this section. For more information, see 4.6, “RAS on the power subsystem” on page 135.

2.6.1 Intelligent Power Distribution Units

By using NVDIMMs for write cache retention, the DS8A00 can help simplify rack power distribution. The DS8A00 systems use 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 2-20 on page 70 shows the iPDU connections.

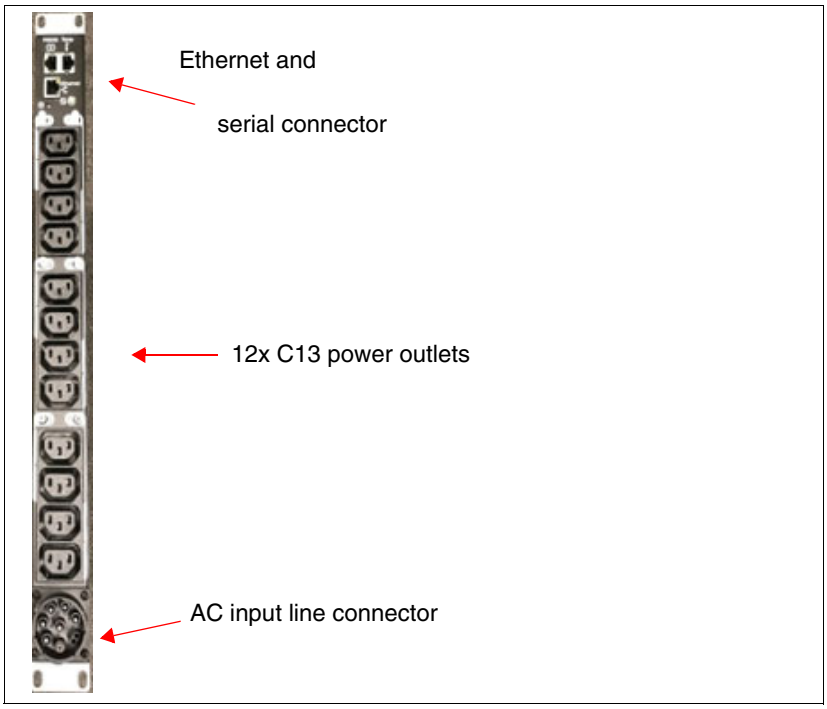


Figure 2-20 iPDU connections

iPDUs are installed in pairs. Each DS8A00 rack has a minimum of one iPDU pair. For models A05, a second pair can be installed in the base frame to provide power for more I/O and storage enclosures.

DS8A00 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 2.6.3, “Backup Power Modules and NVDIMM” on page 73.

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.

DS8A00 HMCs can control system power and monitoring by communicating to the network interfaces of the iPDUs and CPCs.

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 2-21 on page 71 shows redundant Ethernet connections (black and gray network) of the iPDUs and CPCs for a DS8A50, which is equivalent to the network configuration on a DS8A10. The red connection for both rack power control hubs (Ethernet switches) provides backup power from the other rack control hub in the event of local power failure.

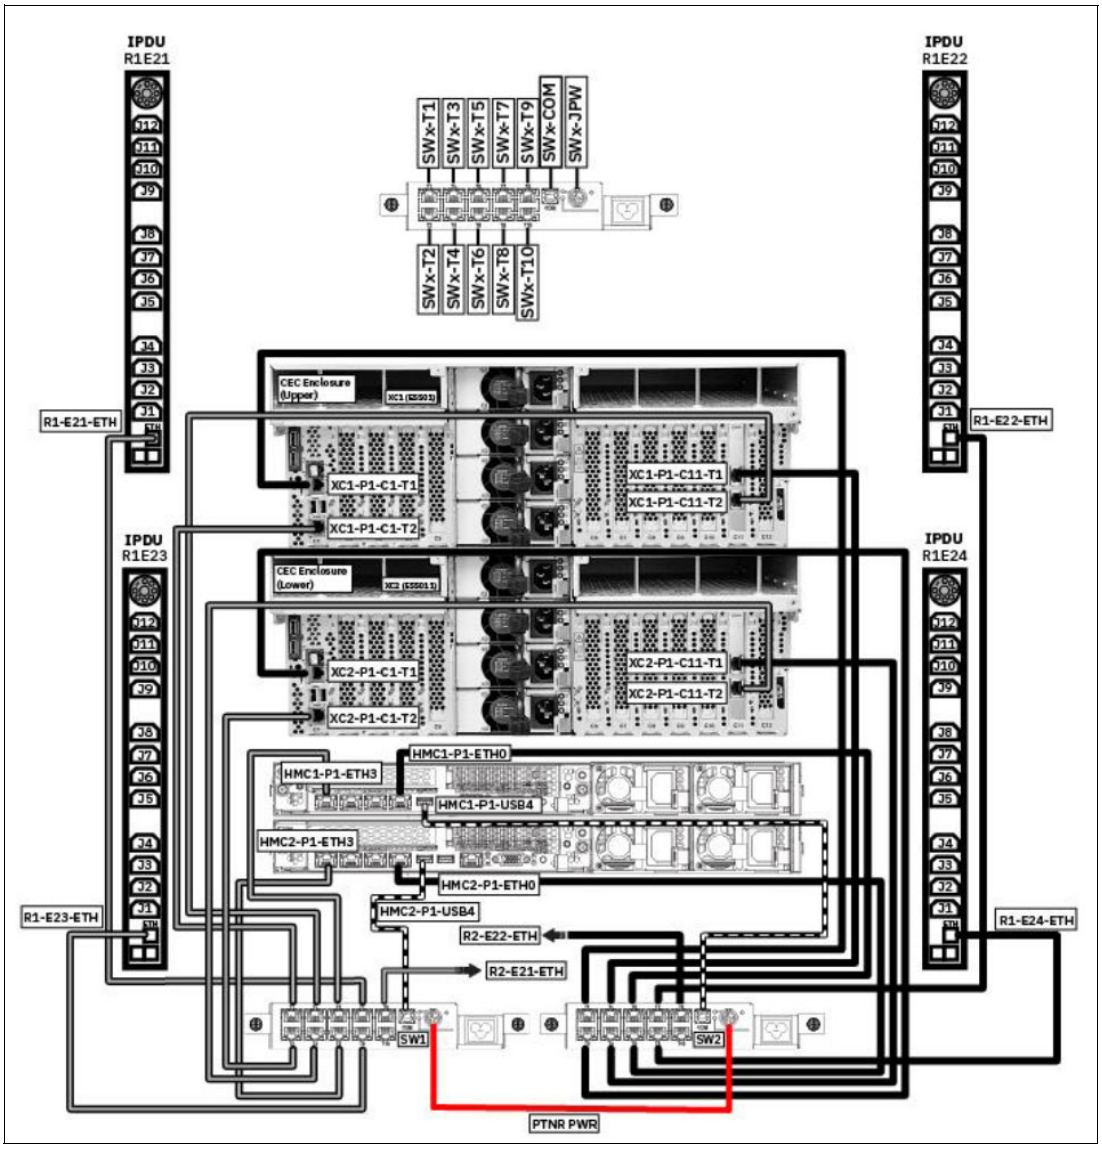


Figure 2-21 DS8A50 Ethernet connections to iPDUs and CPCs

Figure 2-22 shows an example how the power is distributed when two iPDU pairs are installed. Note the power connections of a second HPFE Gen3 storage enclosure or the second I/O enclosure pair.

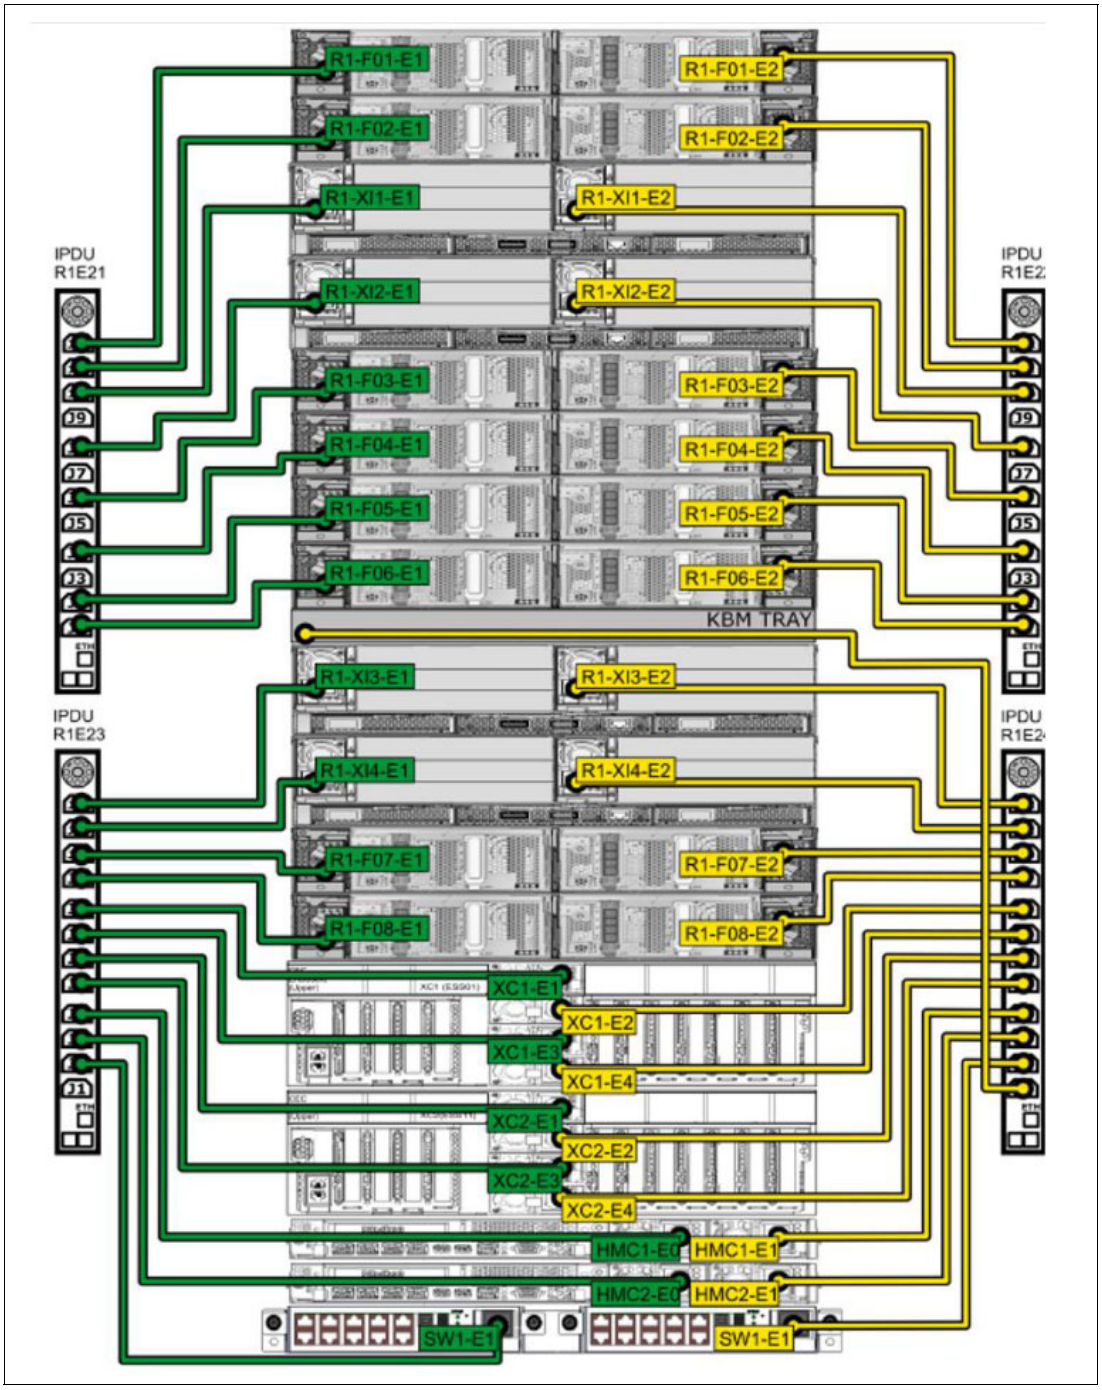


Figure 2-22 DS8A50 iPDU power connections

Adding a model E05 expansion frame to a DS8A80 or DS8A50 system also adds another iPDU pair in that frame, which requires Ethernet connections to the rack control hub in the base frame.

2.6.2 Power domains

To provide redundant power, the iPDUs are organized into groups that are called *power domains*. As shown in Figure 2-22 on page 72, the iPDUs on one side of the racks form a power domain. In 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 Gen3 enclosures, and the HMCs are connected across both power domains. The left power supplies connect to the green domain, and 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.

Important: During data center power maintenance, after the power is restored to the iPDU, validate all PSU input power LEDs before removing power from the partner iPDU. Power supplies in HMCs, I/O enclosures, HPFE Gen3 enclosures, CPC enclosures, rack control hub must have solid green input LEDs before powering off the second iPDU in the pair.

2.6.3 Backup Power Modules and NVDIMM

In DS8A00 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 DS8A00 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 data dump, which was used on earlier DS8000 models to harden NVS data to disk. Figure 2-23 shows a symbolic view of an NVDIMM module.

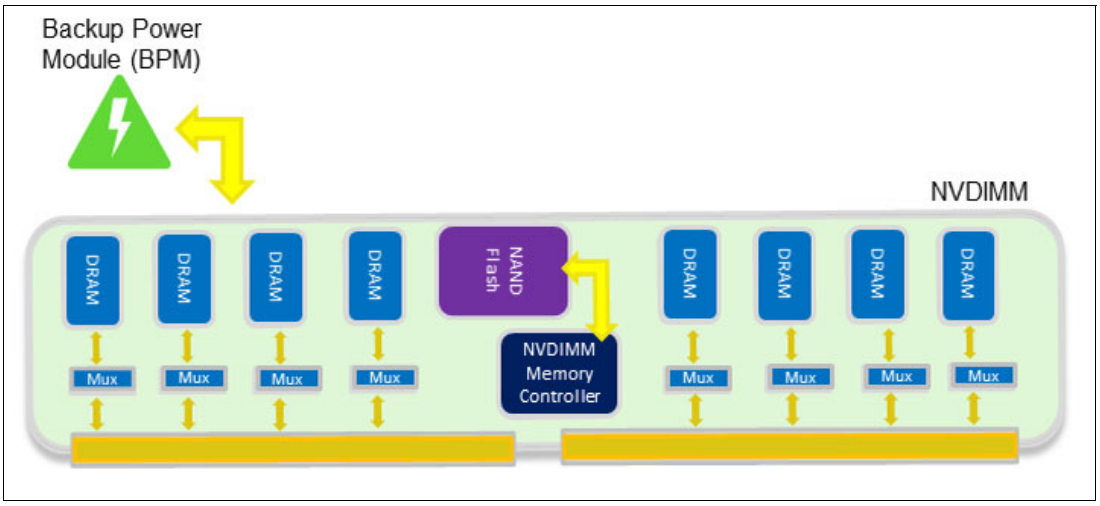


Figure 2-23 NVDIMM logic

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 up to four 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 four BPMs are installed per CPC.

Figure 2-24 shows an example of a BPM.

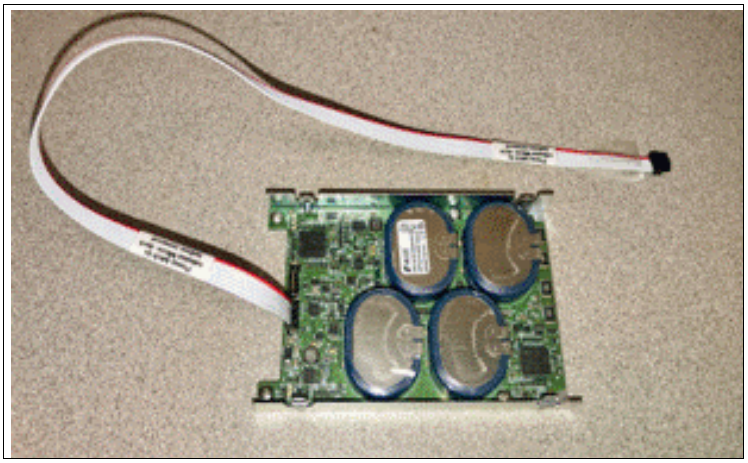


Figure 2-24 Backup Power Module

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 4.2, “CPC failover and failback” on page 116.

NVDIMM configurations use either up to four 16 GB modules or up to four 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.

2.6.4 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 IBM Storage DS8000 10.1 Introduction and Planning Guide, G10-I53-00

Three-Phase, and Single-Phase Power connections

Power attachment to a DS8000 is available by either Three-Phase Power connect (L1, L2, L3, N, PE – Wye, 5 connections; or: L1, L2, L3, PE – Delta, 4 connections), or Single-Phase Power connect (L, N, PE – 3 connections). One of the advantages of three-phase alternating

current is that it delivers almost twice as much power as single-phase systems without the need for twice the number of wires. Also, a better phase compensation is often possible.

Restrictions

For the DS8A00, certain restrictions apply with regards to power attachment:

- FCMs are restricted to up to 2 drive sets (32 drives) per HPFE, if Single-Phase Power is used. This restriction does not apply to industry-standard drives. It also does not apply if 3-phase power is used.
- If you're using the Expansion frame (E05) with Single-Phase Power and have more than 2 HPFE features, you need 4 iPDUs. This restriction does not apply when using 3-phase power attachment.

2.6.5 Enclosure power supply units

The CPCs, I/O enclosures, HMC, and HPFE Gen3 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 DS8A00 is designed for efficient air flow and to be compliant with *hot and cold aisle* data center configurations.

2.7 Hardware Management Console and network

Every DS8A00 base frame is equipped with a pair of Power9 Hardware Management Consoles (HMC), and two private network Ethernet switches also known as rack power control hubs. The secondary HMC is a redundant point of management in the DS8A00.

Figure 2-25 on page 76 shows locations of both HMCs, the keyboard and monitor drawer and the Ethernet switches (rack power control hubs) in a DS8A10 base frame.

The storage administrator runs all DS8A00 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 DS8A00 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 DS8A00 Management Console planning and setup” on page 167.

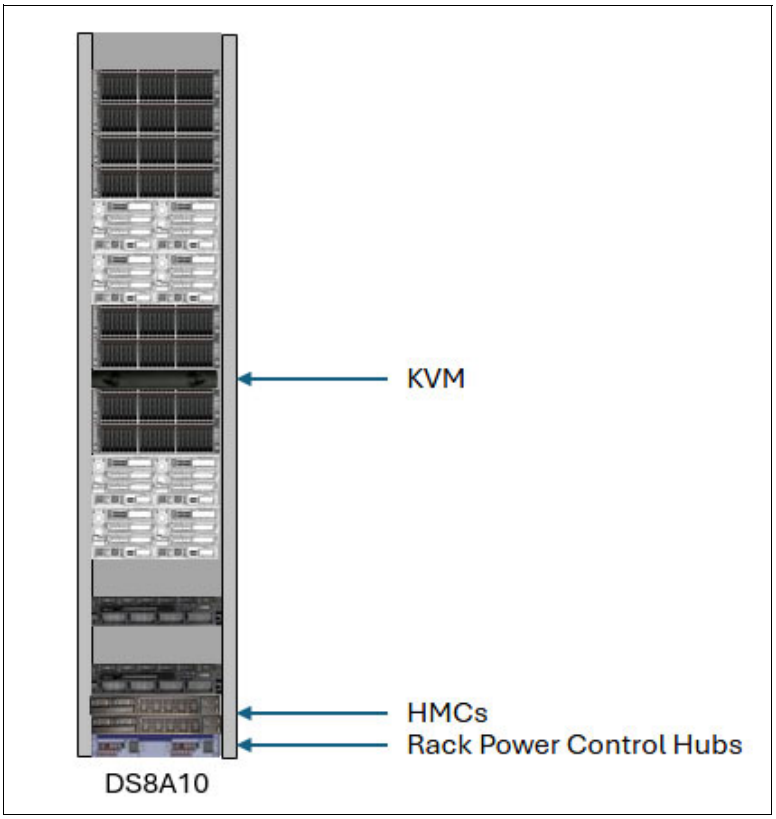


Figure 2-25 DS8A10 location of KVM, HMCs, and Rack Power Control Hubs (Ethernet switches)

Ethernet switches (Rack Power Control Hubs)

The DS8A00 base frame has two 10-port Ethernet switches (Rack Power Control Hubs) installed. 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 switch network, see Figure 2-21 on page 71. The switches receive power from the IPDUs in the Base Frame.

Each HMC also uses two designated Ethernet interfaces for the internal black (eth0) and gray (eth3) networks.

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 E05 is installed for a DS8A80 or DS8A50, an Ethernet connection is established from both Rack Power Control Hubs to the iPDUs in the expansion frame.

Important: The internal Ethernet switches (Rack Power Control Hubs) are for the DS8A00 private networks only. *Do not* connect an external network (or any other equipment) to the black or gray network switches.

3

Virtualization concepts

This chapter describes virtualization concepts as they apply to the IBM DS8A00 system

This chapter covers the following topics:

- “Virtualization definition” on page 78
- “Benefits of virtualization” on page 78
- “Abstraction layers for drive virtualization” on page 79
- “Extent pools” on page 82
- “Terminology for IBM Storage products” on page 106?

© Copyright IBM Corp. 2025. All rights reserved.

3.1 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.

3.2 Benefits of virtualization

Virtualization in IBM DS8A00 includes the following benefits:

- Flexible logical volume configuration:
  - Multiple redundant arrays of independent disks (RAID) types (RAID 6, 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 TiB either with small or large extents (limit of 4 TiB 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.

3.3 Abstraction layers for drive virtualization

Virtualization in the DS8A00 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 if 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 logical volumes.

The basis for virtualization begins with the *physical drives,* which are mounted in storage enclosures and connected to the internal storage servers. DS8A00 uses the High-Performance Flash Enclosure (HPFE) Gen3 storage enclosures exclusively. To learn more about the drive options and their connectivity to the internal *storage servers*, see 2.5, “Flash drive enclosures” on page 67.

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

- Array sites
- Arrays
- Ranks
- Extent pools
- Logical volumes
- LSSs / LCUs

Note: The DS GUI streamlines the process to configure arrays, ranks, and extent pool pairs in a single operation. For more granular control, the DS CLI can be used to manage these entities individually. For more information, see Chapter 10, “IBM DS8A00 Storage Management Command-line Interface” on page 335

3.3.1 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*.

3.3.2 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.

Two RAID levels are supported in the DS8A00:

- RAID 6 (default)
- RAID 10

Each HPFE Gen3 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 no longer supported).

During logical configuration, RAID 6 arrays with the required number of spares are created. Each HPFE Gen3 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+S. Subsequent RAID 6 arrays in the same HPFE Gen3 Pair are 6+P+Q.

Important: Using RAID 6 is recommended, and it is the default in the DS GUI. As drives increase in size, the RAID rebuild times (after one drive failure) also increase. Using RAID 6 reduces the danger of data loss due to a double-drive failure. For more information, see 4.5.1, “RAID configurations” on page 127.

For more information about the sparing algorithm, see 4.5.10, “Spare creation” on page 134.

Figure 3-1 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.

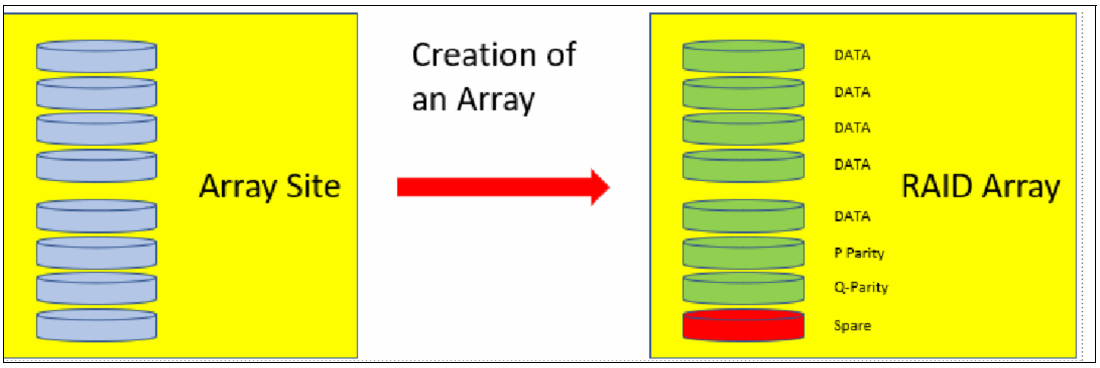


Figure 3-1 Creation of a 5+P+Q+S array

Depending on the selected RAID level and sparing requirements, four types of arrays are possible, as shown in Figure 3-2.

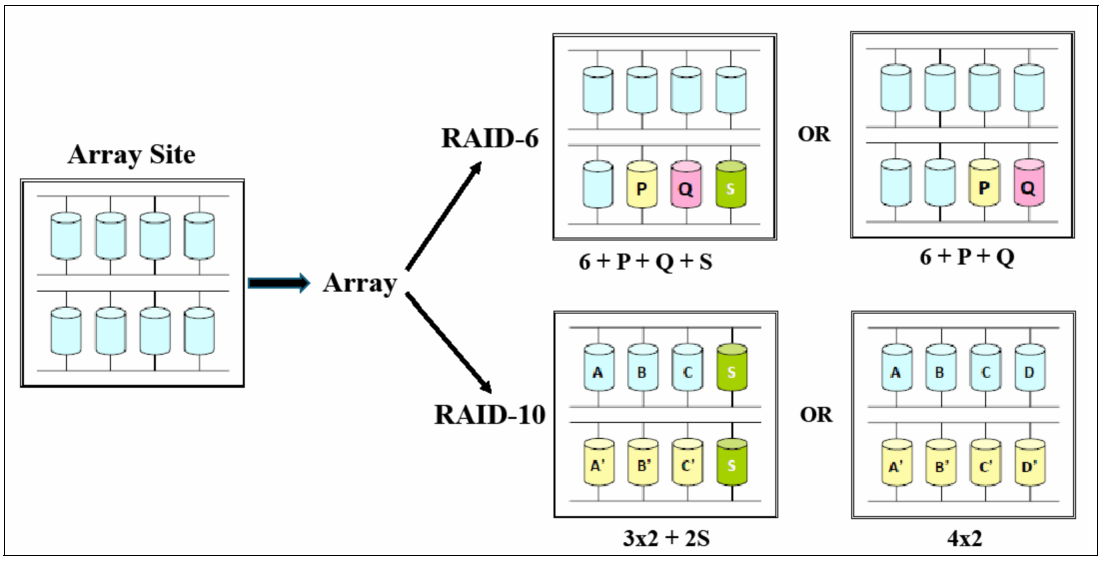


Figure 3-2 Supported RAID array types

Tip: Larger drives have a longer rebuild time. Only RAID 6 can recover from a double drive failure 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 DS8A00.

3.3.3 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 Fixed-Block (FB) or Count key data (CKD) storage types. In the DS8A00, 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 DS8A00 are Full Disk Encryption (FDE)-capable to secure all logical volume data at rest. In the DS8A00, 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 DS8A00 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*, REDP-4500.

Defining ranks

When a new rank is defined, its name is chosen by the DS GUI or 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 3-3 on page 83.
- You can choose between large extents and small extents.

A 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.

In a 3390 Model 1, the large extent size for CKD ranks is 1113 cylinders. The small extent size for CKD ranks 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. 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). However, 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 DS8A00 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 cylinders for small extents. So, 1112 cylinders are wasted if large extents are used.

Note: In the DS8A00 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 3.4.4, “Volume allocation and metadata” on page 96.

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.

Note: The choice of small or large extents has direct correlation with usable capacity, as seen on “System Capacity Limitations” on 2.5, “Flash drive enclosures”.

3.4 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. Historically, extent pools were used to separate drives with different RPMs and capacities. Although you can still use extent pools to separate different flash tiers (Tier 1 and Tier 2), Easy Tier doesn't manage data placement across these 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. In addition, FCM (FlashCore Module) capacity cannot be intermixed with independent storage capacity within 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 the majority of systems, a single pair of extent pools for each rank type (FB or CKD) will provide the best overall performance.

Easy Tier moves data within an extent pool to optimize the placement of the data within the pools.

Note: FCM storage requires Easy Tier to be enabled.

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 DS8A00 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.

In DS8A00, small extents are the default when creating extent pools. If you plan on using large extents for ESE volumes you must create an additional pool pair with large extents. Small and large extents cannot be in the same pool.

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

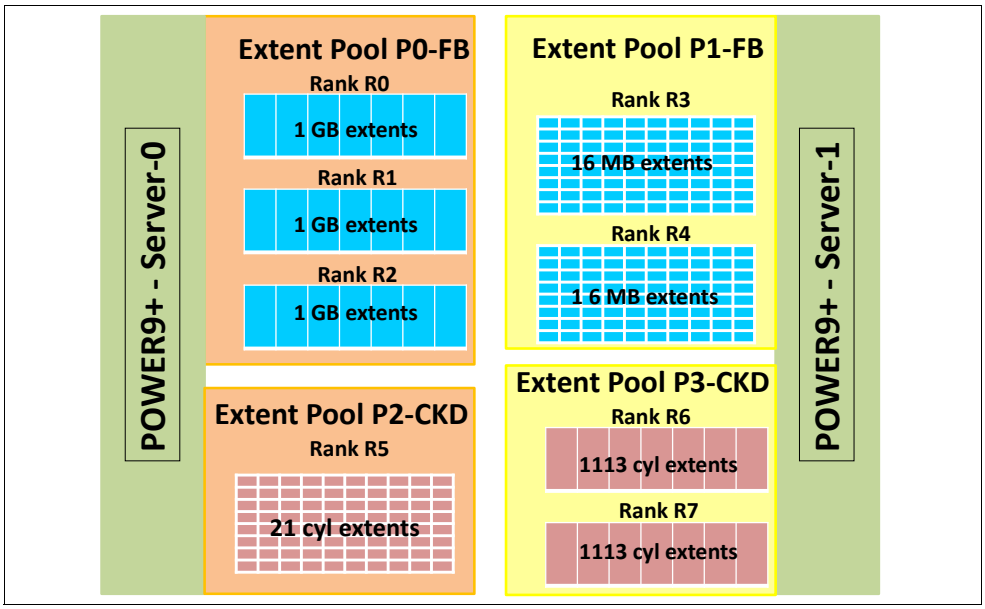


Figure 3-3 Extent pools

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, however in DS8A00, both servers have read/write access to all drives, using all four internal NVMe paths.

Emergency Space Reserve: In DS8000 systems, the *Emergency Space Reserve (ESR)* is a dedicated capacity within an extent pool that activates when the pool reaches a critical threshold. At approximately 98% utilization, the system enters a write-inhibit state, preventing further host writes to avoid overfilling. During this state, the ESR is used to support essential internal operations and maintain metadata integrity, ensuring system stability. This safeguard is designed to prevent catastrophic failures and provide administrators time to take corrective actions, such as adding capacity or migrating volumes.

3.4.1 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 3-4 on page 85. 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 flash drive type (Independent Storage or FlashCore Module)
- Have the same formatting of either Fixed Block (FB) or Count Key Data (CKD)
- Have the same DS8A00 internal server affinity

When merging extent pools, the logical volumes remain accessible to the host systems. Dynamic extent pool merge can be used for the following reasons:

- Consolidation of two extent pools with the equivalent storage type (Independent Storage or FlashCore Module), equivalent formatting (FB or CKD), and equivalent extent size, into one extent pool. If you create a pool that contains more extents, logical volumes can 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.

Note: Migrating volumes from one FCM pool to another FCM pool is not supported with DS8A00 Release 10.1.

- Consolidating Independent Storage 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.

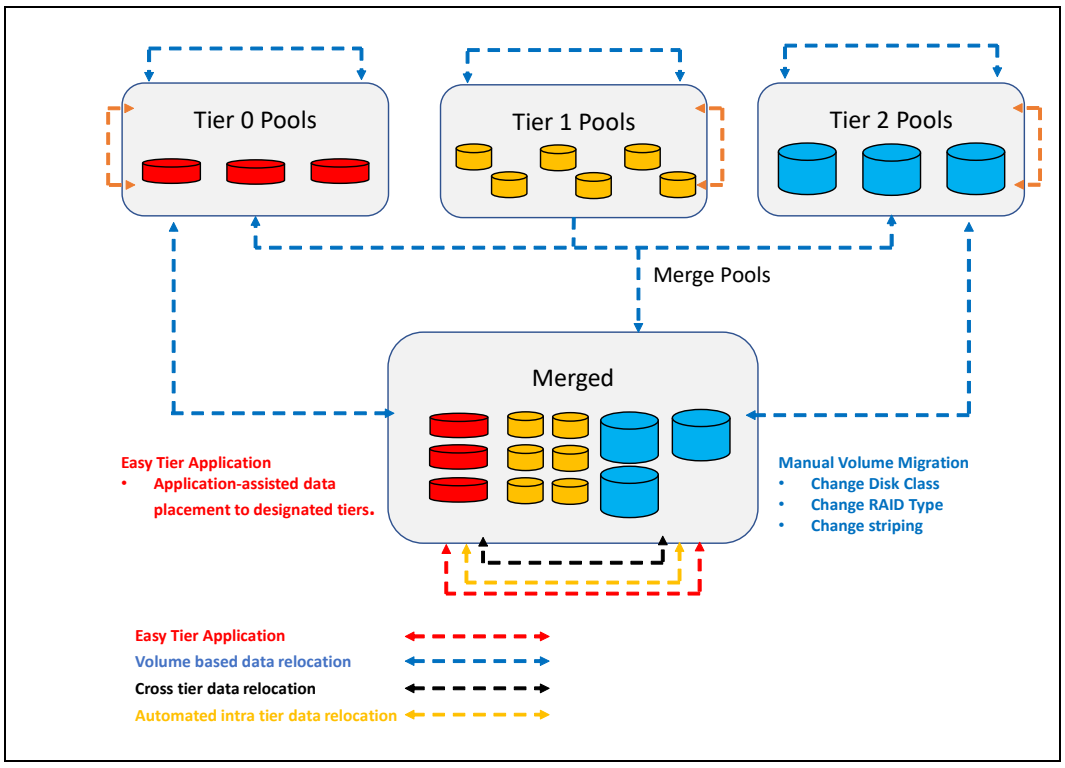


Figure 3-4 Easy Tier migration types

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.

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 (Independent Storage and FlashCore Module)
- If source and target pools have different format types (Fixed Block and Count Key Data).
- 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.

3.4.2 Logical volumes

A *logical volume* consists of a set of extents from one extent pool. The DS8A00 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 x 2^{40} bytes, or 2^{44} bytes) when you use large extents.

Important: DS8000 Copy Services (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 in an IBM SAN Volume Controller (SVC), and is using IBM Spectrum Virtualize CS.

LUNs can be provisioned (allocated) in binary GiB (2^{30} bytes), decimal GB (10^9 bytes), or 512-byte 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 3-5.

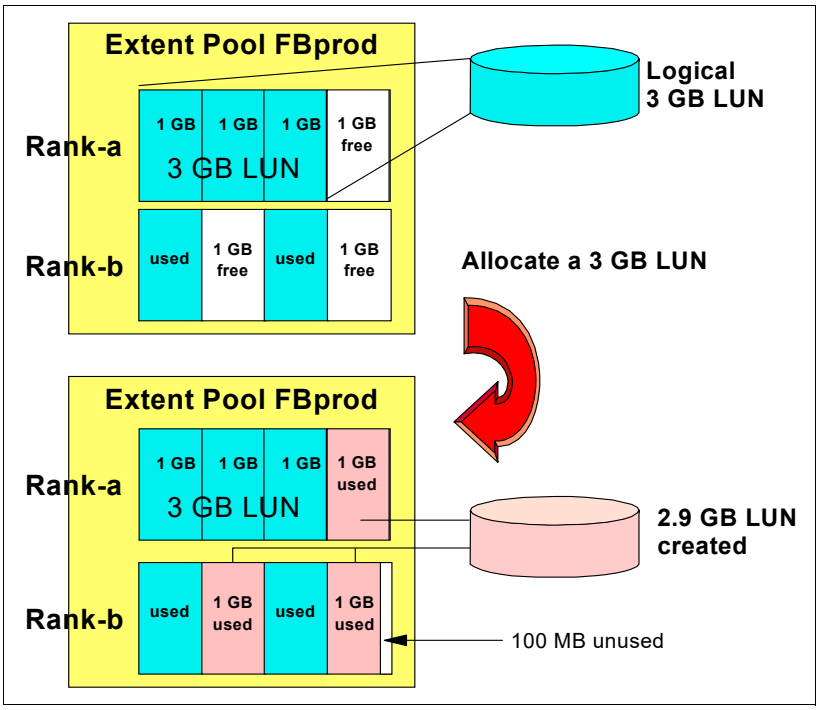


Figure 3-5 Creation of an FB LUN

With small extents, wasted storage is not an issue.

A FB LUN must be managed by a Logical SubSystem (LSS). One LSS can manage up to 256 LUNs. The LSSs are created and managed by the DS8A00, as required. A total of 255 LSSs can be created in the DS8A00.

IBM i logical unit numbers

IBM i LUNs are composed of FB 1 GiB (or 16 MiB) extents and special aspects must be considered with IBM i LUNs. LUNs that are created on a DS8A00 are always RAID-protected. LUNs are based on RAID 6 or RAID 10 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 can use the **`unprotected`** attribute, in which case the DS8A00 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 DS8A00 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 (10^9). Convert these capacities to GiB (2^{30}) when you consider the effective usage of extents that are 1 GiB (2^{30}).

Important: The DS8A00 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 DS8A00 capacity usage for IBM i environments.

The DS8A00 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 3-1. IBM i variable volumes can be dynamically expanded.

Table 3-1 IBM i variable volume sizes

Volume Options		IBM i device size (GB)
Volume	**Protected / Unprotected**	
2107-A01	protected	**8.59**
2107-A02	protected	**17.55**
2107-A04	protected	**70.56**
2107-A05	protected	**35.17**
2107-A06	protected	**141.12**
2107-A07	protected	**282.35**
2107-050	unprotected	**variable**
2107-A81	unprotected	**8.59**
2107-A82	unprotected	**17.55**
2107-A84	unprotected	**70.56**
2107-A85	unprotected	**35.17**

Volume Options		IBM i device size (GB)
Volume	Protected / Unprotected	
2107-A86	unprotected	**141.12**
2107-A87	unprotected	**282.35**
2107-099	protected	**variable**

Example 3-1 demonstrates the creation of both a protected and an unprotected IBM i variable size volume by using the DS CLI.

Example 3-1 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 DS8A00 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 DS8A00 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 DS8A00 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 for AIX on Power Systems servers. For other distributed OSs, read their documentation.

When you create an FB LUN by running the DS CLI command **`mkfbvol`**, add the option **`-t10dif`**. If you query a LUN with the **`showfbvol`** command, the data type is FB 512T instead of the standard FB 512 type.

Important: Because the DS8A00 internally always uses 520-byte sectors to support IBM i volumes, no 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 (count key data) 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, a Logical Control Unit (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 3.4.5, “Logical subsystems” on page 100.

On a DS8A00, 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 can also stripe a volume across the ranks. For more information, see “Storage pool striping: Extent rotation” on page 92.

Figure 3-6 shows an example of how a logical volume is provisioned (allocated) with a CKD volume.

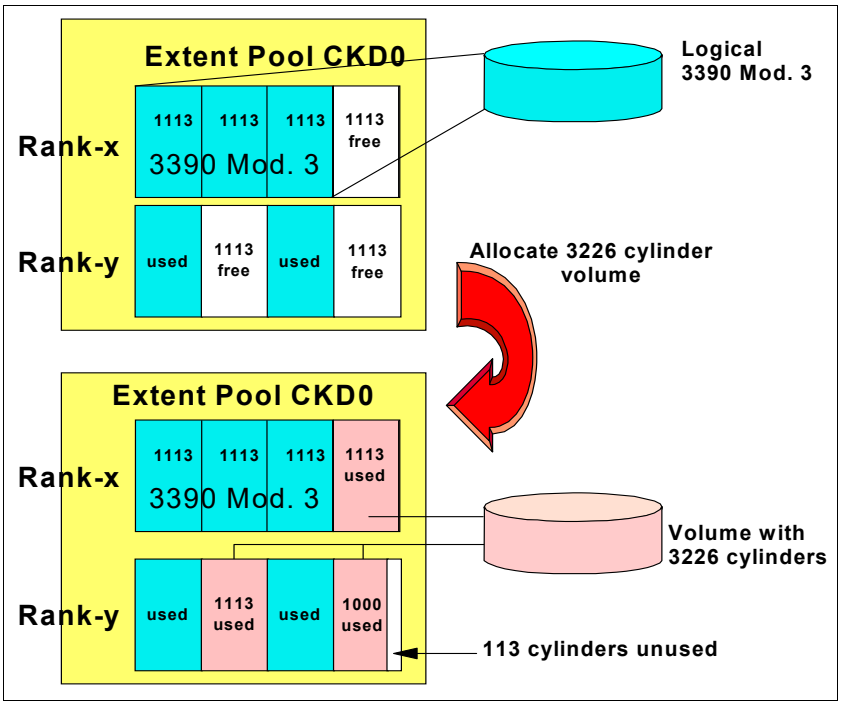


Figure 3-6 Allocation of a Count Key Data logical volume

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 DS8A00 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 DS8A00 internal server and have the same paths to the DS8A00. SMF can provide reports at the Alias Management Group level.

Figure 3-7 illustrates how SuperPAV works.

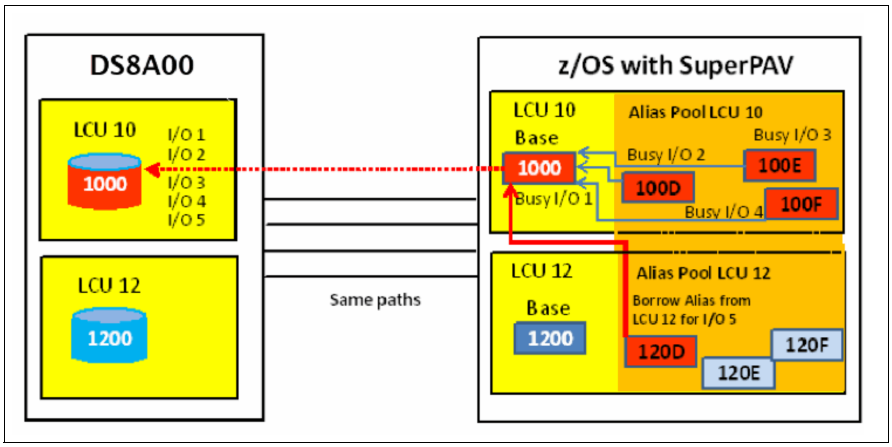


Figure 3-7 SuperPAV

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 `IECIOSxx` 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 3-2).

Example 3-2 Display 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.

Note: For more information regarding HyperPAV and SuperPAV, refer to*IBM DS8000 and IBM Z Synergy (DS8000 Release 9.3 and IBM z/OS 2.5)*, REDP-5186.

3.4.3 Allocation, deletion, and modification of 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 DS8A00 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 DS8A00 maintains a sequence of ranks. The first rank in the list is randomly picked at each power-on of the storage system. The DS8A00 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 3-8.

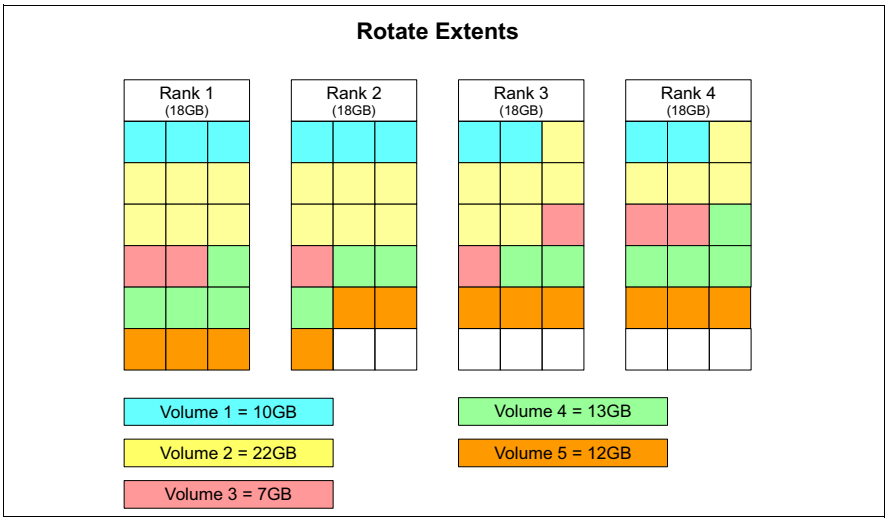


Figure 3-8 Rotate extents

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 3-9.

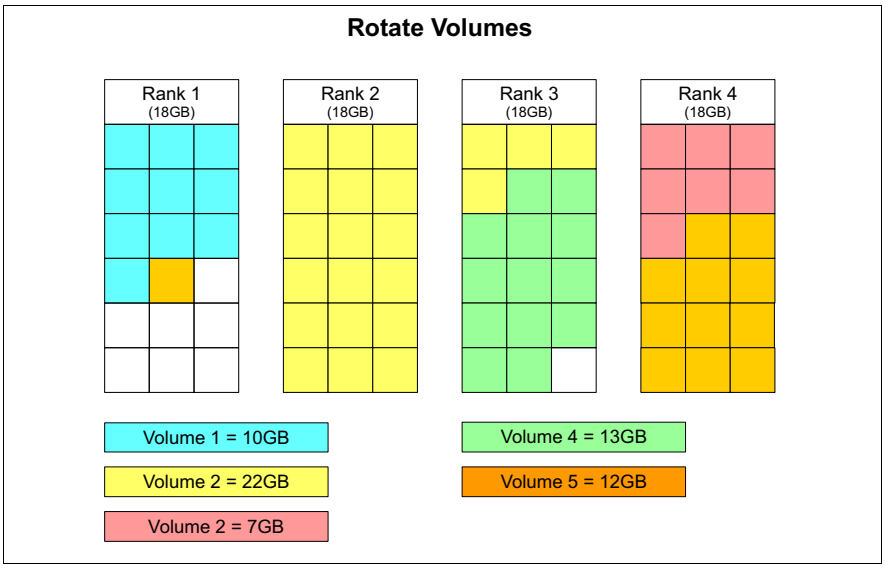


Figure 3-9 Rotate volumes

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`.

Easy Tier's default allocation order prioritizes the highest-performance drive classes. In the case of a DS8A00, this typically means Flash Tier 1 followed by Flash Tier 2. The system allocates new data to the fastest available tier to optimize performance.

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 DS8A00 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 OS 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

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 DS8A00 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 drive type, same extent size, and same formatting.* However, DVR is not supported between two FCM pools in R10.1.

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 any time 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.

3.4.4 Volume allocation and metadata

In the DS8A00 and DS8900F, 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 DS8A00:

- 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. Each large metadata extent contains 64 metadata subextents for FB and 53 metadata subextents 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 3-10 on page 97 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.

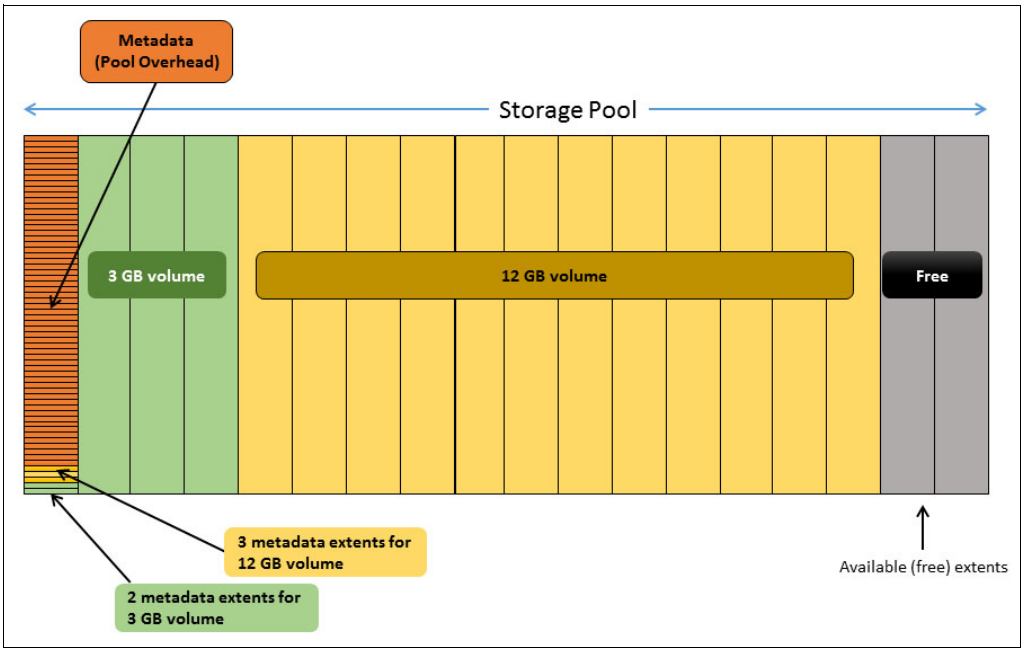


Figure 3-10 Metadata allocation

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 × 2 + total volume extents/10)/64 and rounded up to the nearest integer
- CKD pool overhead = (number of volumes × 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 x 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 x 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 3.5, “Terminology for IBM Storage products” on page 106.

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 **`ETSRMode`** 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.

Note: Volumes configured using FlashCore Module (FCM) drives must be thin 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 DS8A00 runs low on raw capacity, more is added. This approach also assumes that the DS8A00 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. The DS8A00 creates thin-provisioned volumes by default.

CKD thin provisioning

Thin provisioning is also possible with CKD volumes. See *IBM DS8000 Thin Provisioning (Updated for Release 9.3)*, 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.

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 percentage of available capacity is greater than the extent pool / repository threshold.
- `1`: The percentage of available capacity is greater than zero but less than or equal to the extent pool / repository threshold.
- `10`: The percentage 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.

Note: Trap 223 is used for reporting Extent Pool Capacity warnings. See 12.2.3, “Thin-provisioning SNMP” on page 431for examples.

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 3-3 provides an example of creating and modifying an extent pool with a defined operational limit that cannot be exceeded.

Example 3-3 The opratiolimit parameter

```
dscli> mkextpool -rankgrp 0 -stgtype fb -opratiolimit 3.5 -encryptgrp 1 my_pool
CMUC00000I mkextpool: Extent pool P3 successfully created.
dscli> 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 DS8000 Thin Provisioning (Updated for Release 9.3)*, REDP-5343.

3.4.5 Logical subsystems

A LSS is another logical construct. It can also be referred to as a 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 DS8A00. Each of them have an identifier of 00–FF. 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 DS8A00 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 102.

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 DS8A00 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 DS8A00. The LUN is sent, so it can include the LUN in the “frame” that is sent to the DS8A00 when it wants to run an I/O operation on the volume. This method is how the DS8A00 knows which volume 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 DS8A00. 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 DS8A00 when it wants to run an I/O operation on the volume. This method is how the DS8A00 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 DS8A00, the adapter checks the LSS or LCU identifier to know which internal server to pass the request to inside the DS8A00. For more information about host access to volumes, see 3.4.6, “Volume access” on page 102.

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 3-11.

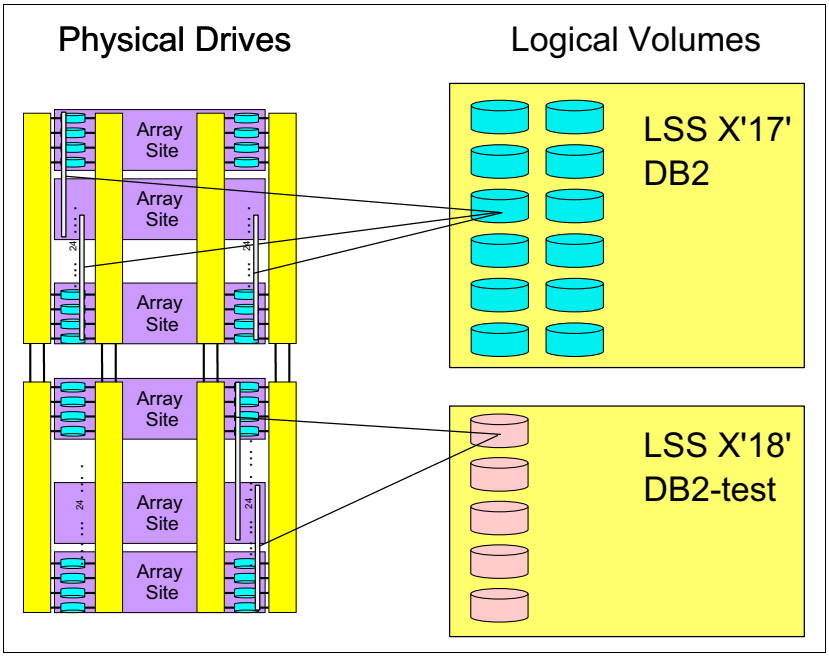


Figure 3-11 Grouping of volumes in LSSs

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 3-12 shows the concept of LSSs and address groups.

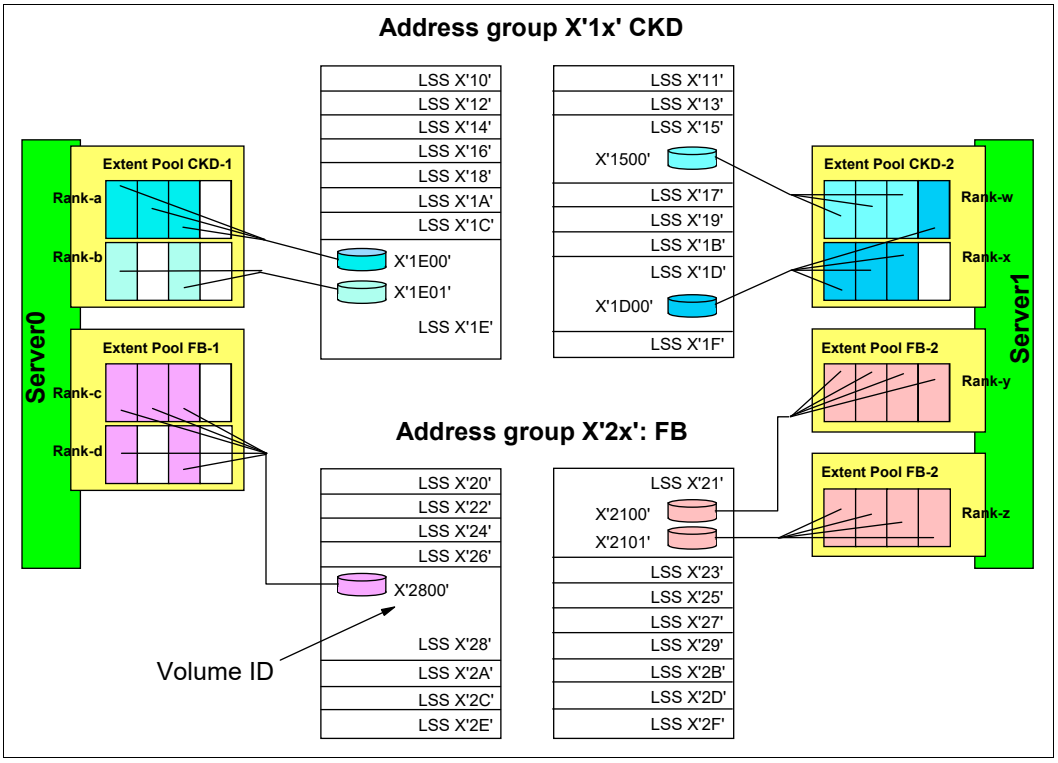


Figure 3-12 Logical storage subsystems

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 3-12 shows only one address group that is used with each extent pool.

3.4.6 Volume access

The DS8A00 provides mechanisms to control host access to LUNs. Usually, a host system features two or more HBAs and the system needs access to a group of LUNs. For management of system access to logical volumes, the DS8A00 introduces the concept of host attachments and volume groups.

Host attachment

HBAs are identified to the DS8A00 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 the 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 DS8A00 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 DS8A00 is 8,192. This host definition is required only for open systems hosts. Any IBM Z servers can access any CKD volume in a DS8A00 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 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. 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. 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 DS8A00 supports a maximum of 8,320 volume groups.

Figure 3-13 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 3-13. 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`.

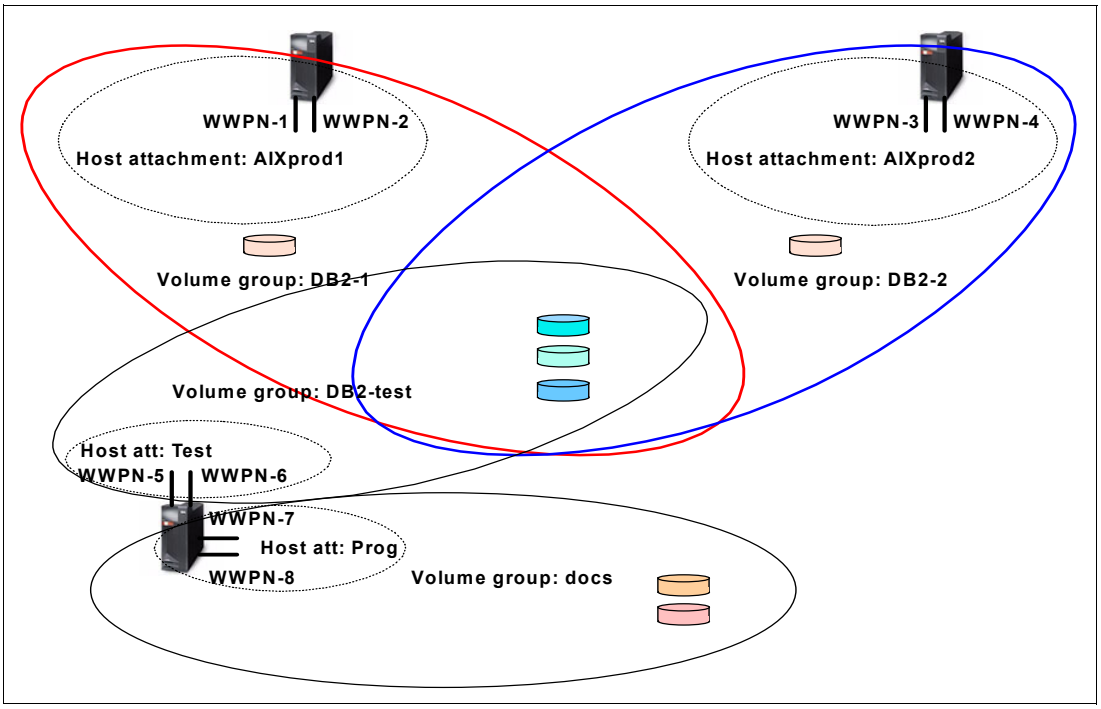


Figure 3-13 Host attachments and volume groups

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.

3.4.7 Virtualization hierarchy summary

As you view the virtualization hierarchy (as shown in Figure 3-14 on page 105), 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 3-14 shows the virtualization hierarchy.

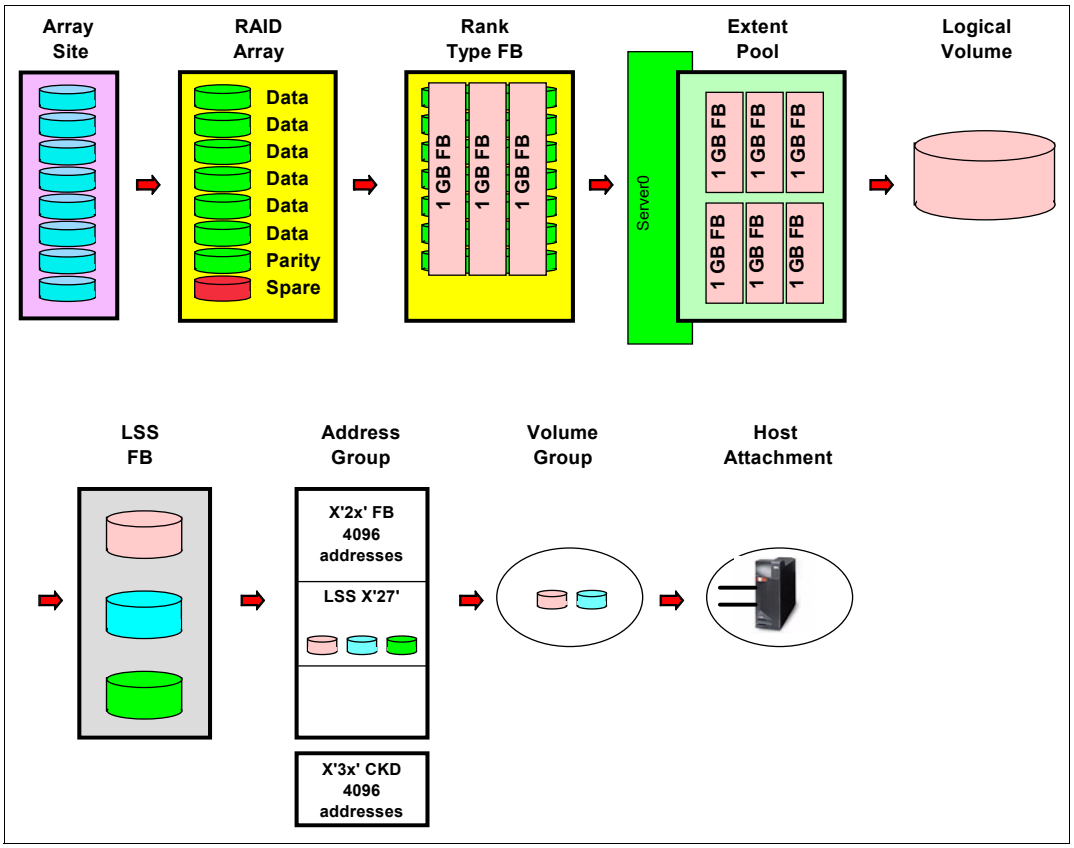


Figure 3-14 Virtualization hierarchy

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.

3.5 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 (IS)

  The reported capacity of the drives in the system before RAID and formatting are applied.

- Total logical capacity (FCM)

  The total amount of capacity that can be written (also known as total effective capacity).

- Usable capacity (IS)

  The amount of capacity that is available for storing data on a system, pool, or array after formatting and RAID techniques are applied.

- Total Physical Capacity (FCM)

  The amount of raw capacity that can be stored on media after data reduction.

- Used Logical Capacity (FCM)

  The total amount of capacity written from the host before data reduction techniques are applied.

- Available (usable) Capacity (IS)

  The amount of usable capacity that is not yet used in a system, pool, or array.

- Used Physical Capacity (FCM)

  The amount of usable capacity that is taken up by data after data reduction techniques are applied.

- 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.

  Provisioned capacity was previously called *virtual capacity* in this chapter.

- Written Capacity

  The amount of usable capacity that is written in a pool or system.

- 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.

4

IBM DS8A00 reliability, availability, and serviceability

This chapter describes the reliability, availability, and serviceability (RAS) characteristics of the IBM DS8000 (10th Generation). The 10th Generation of DS8000s is referred to as the DS8A00.

This chapter covers the following topics:

- “DS8A00 processor complex features” on page 110
- “CPC failover and failback” on page 116
- “Data flow in the DS8A00” on page 120
- “RAS on the Hardware Management Console” on page 124
- “RAS on the storage system” on page 126
- “RAS on the power subsystem” on page 135
- “Other features” on page 138

© Copyright IBM Corp. 2025. All rights reserved.

4.1 DS8A00 processor complex features

RAS is a key principle in the design of the DS8A00. A combination of hardware features, software capabilities, design considerations, and operational guidelines work together to ensure the DS8A00 delivers high reliability. At the heart of the DS8A00 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 Featuring PCIe Gen 4 Technology*, REDP-5595.

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.

4.1.1 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.
- 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.

4.1.2 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 DS8A00. You can read more about the Power9 and processor configuration from the DS8A00 architecture point of view in 2.3.1, “IBM Power9 processor-based CPCs” on page 55.

Power9+ RAS features

The following sections describe the RAS 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 DS8A00 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.

4.1.3 Cross-cluster communication

In the DS8A00, the I/O enclosures are connected point-to-point, and each CPC uses a PCIe architecture. DS8A00 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 tiny portion of the overall available PCIe bandwidth, so XC communication traffic has a negligible effect on I/O performance.

Figure 4-1 shows the redundant PCIe fabric design for XC communication in the DS8A00 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.

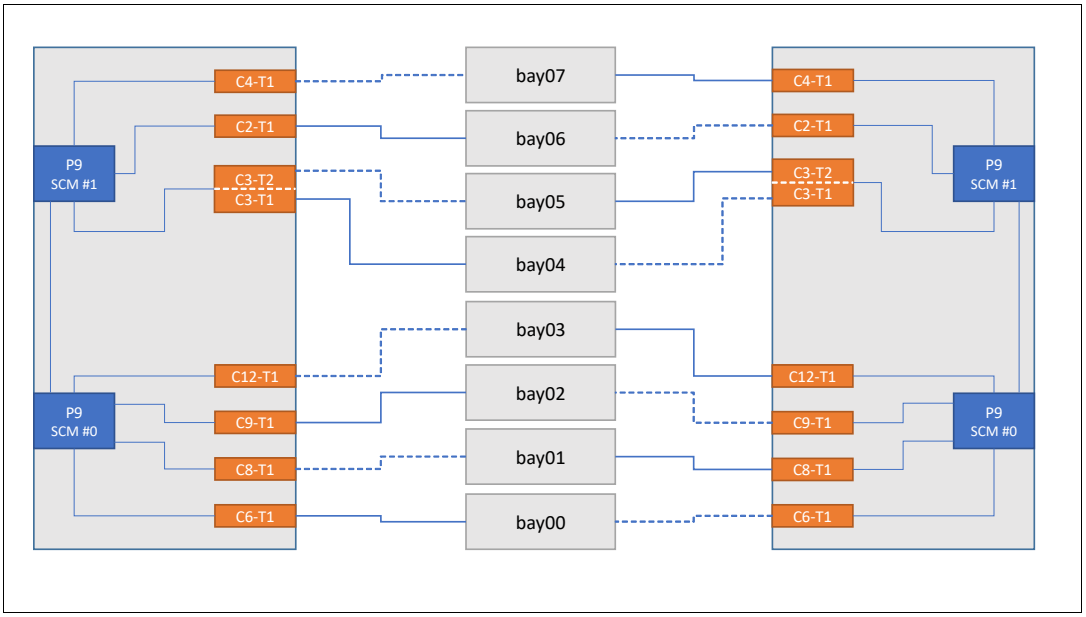


Figure 4-1 DS8A00 XC communication through the PCIe fabric and I/O enclosures

4.1.4 Environmental monitoring

The environment (power, fans, and temperature) is monitored by the flexible service processor (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 4-1.

Example 4-1 The showsu command

```
dscli> showsu
Date/Time: October 9, 2025 12:25:43 PM EDT IBM DSCLI Version: 7.10.10.348 DS: -
Name            -
desc            -
ID              IBM.2107-13xxxx0
Model           A08
WWNN            5005076303FFF854
config          Undefined
pw state        On
pw mode         Remote Manual
reqpm           Remote Manual
System Memory   4864.0 GB
MTS             IBM.5341-13xxxx0
ER Test Mode    Disabled
ER Recorded     2025-10-09T12:24:05-0400
ER Power Usage 5416
ER Inlet Temp   -
ER I/O Usage    18
ER Data Usage   616
```

4.1.5 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.

4.2 CPC failover and failback

To understand the process of CPC failover and failback, review the logical architecture of the DS8A00. For more information, see Chapter 3, “Virtualization concepts” on page 77.

4.2.1 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.

All DS8A00 configurations use 1/16th of system memory except for the largest models, A08 and A05: DS8A80 Model A08, which uses 256 GB of NVS for its 4.8 TB of total system memory and DS8A50 Model A05 with 3.5 TB of total system memory, which uses 192 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 DS8A00 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 4-2 on page 117, 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 alternate 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 alternate 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 alternate node.

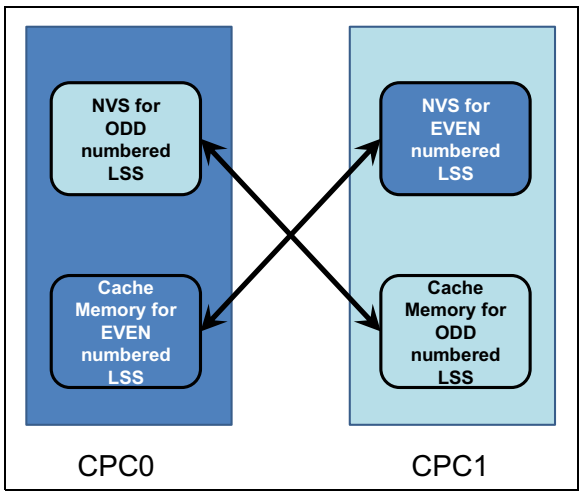


Figure 4-2 NVS write data when both CPCs are operational

Under normal operation, both DS8A00 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.

4.2.2 Failover

In the example that is shown in Figure 4-3, CPC0 failed. CPC1 must take over all of the CPC0 functions. All storage arrays are accessible by both CPCs.

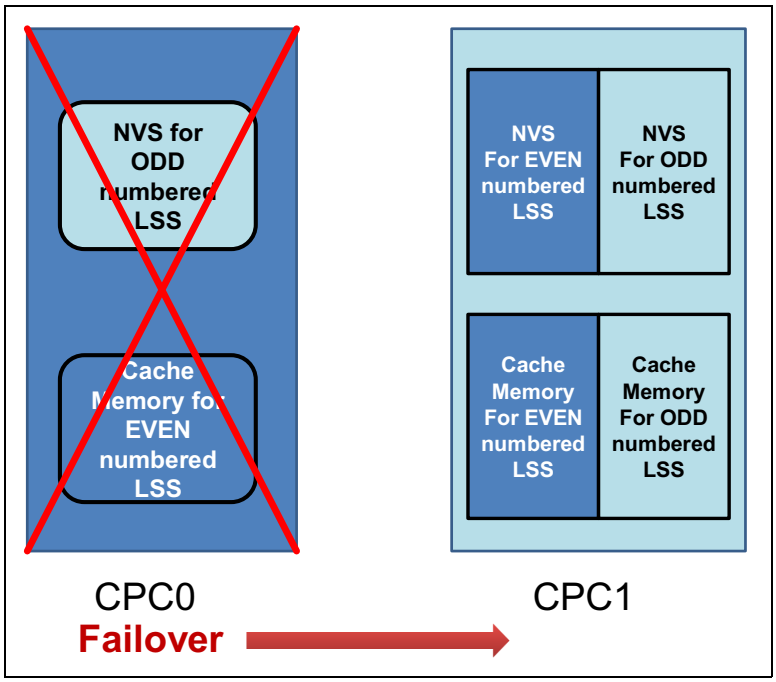


Figure 4-3 CPC0 failover to CPC1

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 DS8A00 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 DS8A00 operates as shown in Figure 4-3 on page 117. 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 DS8A00 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 DS8A00 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.

4.2.3 Failback

The failback process is initiated when the DS8A00 determines that the failed CPC has been repaired and is ready to resume operations. The DS8A00 then automatically starts the resume action, or if manual intervention is required, an IBM SSR or remote support engineer will initiate the failback process.

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 DS8A00 do not affect I/O operation latency by more than 7 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.

4.2.4 NVS and power outages

DS8A00 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 E05 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 DS8A00 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 DS8A00 stores the NVS contents on non-volatile DIMMs (NVDIMMs). Each CPC contains up to four 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: DS8A00 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 DS8A00 must be powered on manually from HMC 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 DS8A00 is powered on automatically.

For more information about how to set power control on the DS8A00 system, see IBM Documentation.

After the DS8A00 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.

4.3 Data flow in the DS8A00

The DS8A00 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 60.

4.3.1 I/O enclosures

As shown in Figure 4-1 on page 114, each CPC on a DS8A80 and DS8A50 can be connected to up to four I/O enclosures within the base frame, and up to eight I/O enclosures when the expansion frame is installed. Each I/O enclosure functions as an extension of the CPCs. Both DS8A10 models (A00 and A10) have a maximum of four I/O enclosures.

The DS8A00 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 four double-stacked enclosure cooling fans. The PSUs and enclosure fans can be replaced concurrently without disruption to the I/O enclosure.

4.3.2 Host connections

Each DS8A00 32 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 4-4 on page 121 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 DS8A00, as shown in Figure 4-4 on page 121, 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. Figure 4-4 on page 121 shows a single-path host connection.

Important: For host connectivity, hosts that access the DS8A00 must have at least two connections to I/O ports on separate host adapters in separate I/O enclosures.

A more robust design is shown in Figure 4-5 on page 121, 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.

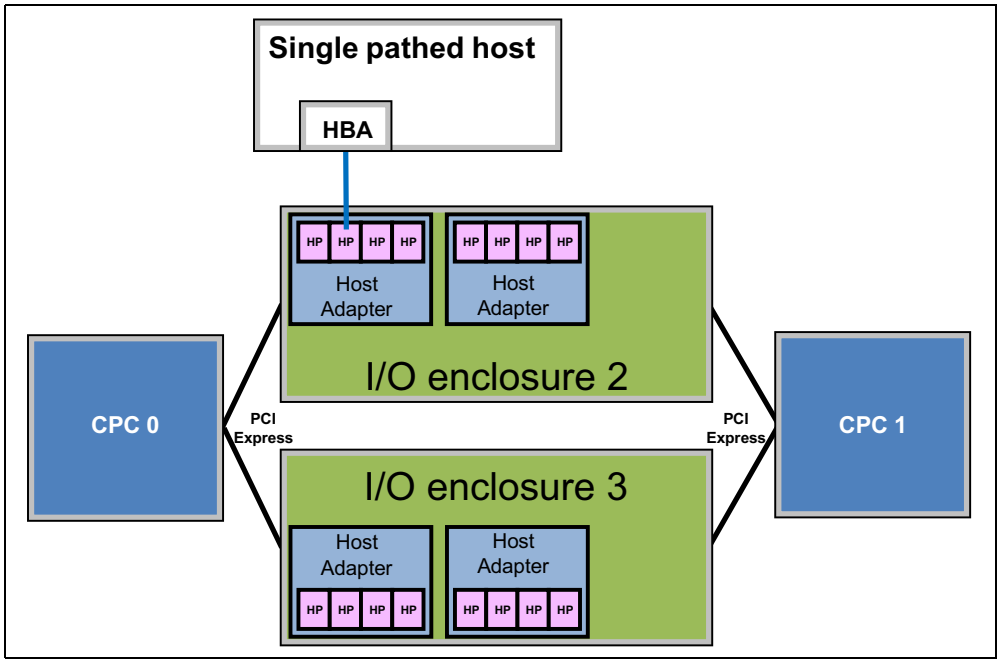


Figure 4-4 A single-path host connection

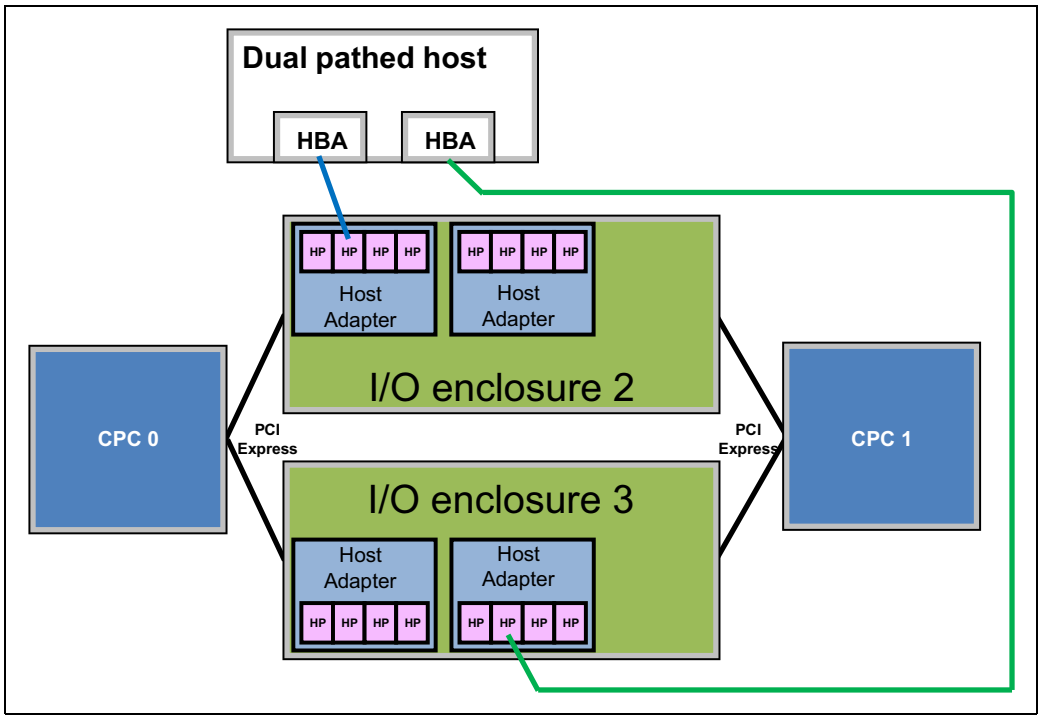


Figure 4-5 A dual-path host connection

SAN and FICON switches

Because many hosts can connect to the DS8A00 by using multiple paths, the number of host adapter ports that are available in the DS8A00 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 DS8A00. 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 DS8A00 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 DS8A00 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 DS8A00, see “T10 Data Integrity Field support” on page 88.

Robust and resilient SAN data transfer

Forward Error Correction (FEC) is enabled on the host adapters to intercept and correct bit errors for utmost reliability. FEC is also implemented on IBM 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 DS8A00. 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 in the past 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 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 DS8A00 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 DS8A00 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, as well as by the DS8A00 storage system.

A physical FICON path is established when the DS8A00 port sees light on the fiber, for example, a cable is plugged in to a DS8A00 host adapter, a processor or the DS8A00 is powered on, or a path is configured online by z/OS. Logical paths are established through the port between the host, and part of or all of the LCUs in the DS8A00 are controlled by the HCD definition for that host. This configuration happens for each physical path between an IBM Z host and the DS8A00. Multiple system images can be in a CPU. Logical paths are established for each system image. The DS8A00 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 DS8A00 service actions.

CUIR is available for the DS8A00 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 DS8A00 service actions.

CUIR also allows the DS8A00 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 DS8A00 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 DS8A00.

4.3.3 Metadata checks

When application data enters the DS8A00 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 DS8A00. 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 DS8A00 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 DS8000 T10 DIF support for FB volumes. For more information about T10 DIF implementation in the DS8000, see “T10 Data Integrity Field support” on page 88.

4.4 RAS on the Hardware Management Console

The Hardware Management Console (HMC) is used to configure, manage, and maintain the DS8A00. Two Power9 HMCs (the *primary* and the *secondary*) are included in every DS8A00 base frame. The DS8A00 HMCs work with IPv4, IPv6, or a combination of both IP standards. For more information about the HMC and network connections, see 5.3, “Network connectivity planning” on page 159.

The HMC is the DS8A00 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), DS 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.

4.4.1 Licensed Internal Code updates

The DS8A00 contains many discrete redundant components. The DS8A00 architecture provides concurrent code updates. This ability is achieved by using the redundant design of the DS8A00. The following components have firmware that can be updated concurrently:

- FSP and IBM Power firmware
- iPDU
- Host adapters
- Flash enclosure
- Device adapters (DAs)
- Flash drives
- I/O enclosure

The DS8A00 CPCs have an OS (AIX) and Licensed Machine Code (LMC) that can be updated. As IBM continues to develop and improve the DS8A00, new releases of firmware and LMC are becoming available that offer improvements in function and reliability. For more information about LIC updates, see Chapter 11, “Licensed Machine Code” on page 405.

4.4.2 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 DS8A00 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 4-2.

Example 4-2 Typical email notification output

```
REPORTING SF MTMS:   2107-A08*130xxx0
FAILING SF MTMS:   2107-A08*130xxx0
REPORTING SF LPAR:   unknown
PROBLEM NUMBER:      491
PROBLEM TIMESTAMP:   Oct 11, 2025 8:17:07 AM MST
REFERENCE CODE:      BEF00012
************************  START OF NOTE LOG  ************************** BASE RACK
ORDERED MTMS 5341-A08*1301520 LOCAL HMC MTMS 7063CR2*7897C4D HMC ROLE Primary
LOCAL HMC OUTBOUND CONFIG SSL only     FTP: disabled
REMOTE HMC MTMS 7063CR2*78986ED HMC ROLE Secondary
REMOTE HMC OUTBOUND CONFIG Internet only     FTP: disabled
AOS STATUS Not Configured AOS VERSION
AOS ACL=(DS8k, Storage) AOS TRACE Enable RSC STATUS Running RSC PASSWORD Not
Required HMC CE default  HMC REMOTE default HMC PE default  HMC DEVELOPER default
2107 BUNDLE 10.11.7.0
HMC BUILD 2509191304
LMC LEVEL v25.101.1 build level 20251007.1 FIRMWARE LEVEL SRV0 01VL950193  SRV1
01VL950193
STORAGE FACILITY 2107-A08*130xxx0,
PARTITION HOST NAME
PARTITION STATUS SFI unknown SVR 9009-42G*7897001 LPAR unknown STATE = FIRST
REPORTED TIME Oct 11, 2025 8:17:07 AM MST LAST REPORTED TIME Oct 11, 2025 8:17:07
AM MST CALL HOME RETRY #0 of 12 on Oct 11, 2025 8:17:10 AM MST.
```

```
REFCODE BEF00012
AUTO EED Disabled   AUTO PE Package Disabled
SERVICEABLE EVENT TEXT
HMC=7063CR2*7897C4D: HeartBeat. This is an informational message on
2107 for sending home data. This ServiceableEvent record should be closed.
No Frus to display.
*************************  END OF NOTE LOG  *************************
```

For more information about planning the connections that are needed for HMC installations, see Chapter 6, “IBM DS8A00 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 DS8A00. 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.

4.5 RAS on the storage system

The DS8A00 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 DS8A00 system supports RAID 6 and RAID 10.

Note: Due to the added resiliency of RAID 6, RAID 5 is no longer supported on DS8A00 Models.

4.5.1 RAID configurations

The following RAID configurations are supported on DS8A00:

- 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.

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 employs 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 (StorM) is an easy-to-use web tool that is available 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:

- Within one High-Performance Flash Enclosure (HPFE) Gen3 pair of six array sites, a RAID intermix is allowed, but no intermix of FCM drives with Flash Tier 1 or Flash Tier 2 drives is supported.

For the latest information about supported RAID configurations or to request an RPQ / SCORE, contact your IBM Technical Sales.

4.5.2 Drive path redundancy

Each flash drive inside an HPFE in the DS8A00 is attached to two PCIe Gen4 NVMe switch expanders, which are contained within the Enclosure Service Modules (ESMs) in the flash drive enclosure.

Each flash drive has two separate connections to the enclosure backplane. This configuration allows a flash drive to be simultaneously attached to both NVMe expander switches. If either ESM is removed from the enclosure, the NVMe 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.

For more information about the drive subsystem of the DS8A00, see 2.5, “Flash drive enclosures” on page 67.

4.5.3 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.

The DS8A00 introduces a dual-active design that allows each node to connect directly to both adapters in a pair. This setup boosts performance and fault tolerance by ensuring that if one adapter fails, I/O can be redirected locally without needing to switch arrays or change roles. This design not only reduces disruptions during software updates but also enhances how adapter failures are managed, especially in single CEC mode.

Figure 4-6 shows the connections for the DA and flash enclosure pair.

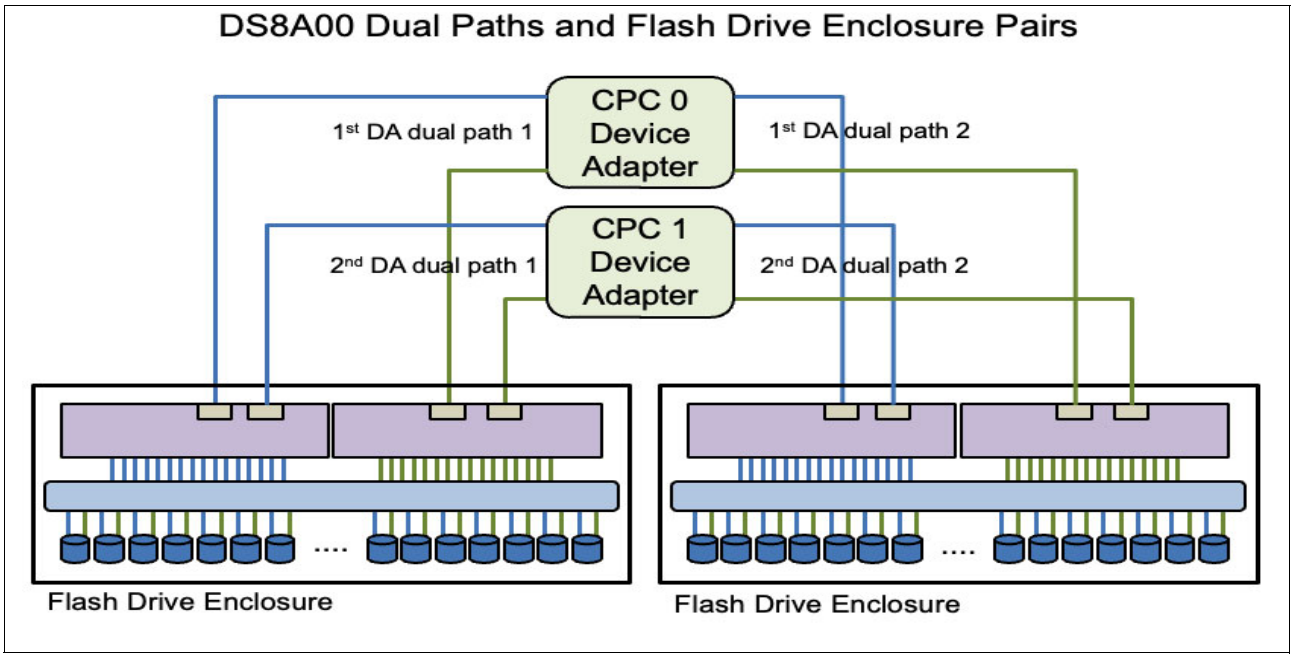


Figure 4-6 Dual paths on a flash drive enclosure

4.5.4 Predictive failure analysis

The flash drives that are used in the DS8A00 incorporate Predictive failure analysis (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 4.5.9, “Smart Rebuild” on page 133. This copy ability avoids using RAID recovery to reconstruct all the data on to the spare drive.

4.5.5 Disk scrubbing

The DS8A00 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 increase to a level beyond error correction capabilities.

4.5.6 RAID support

Arrays can be configured as RAID 6 or RAID 10.

Important: RAID 6 is now the default and preferred setting for the DS8A00. RAID 10 continues to be an option for Flash Tier 1, Flash Tier 2, and FCM drives.

The DS8A00 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 Gen3 pair in a DS8A00 can contain up to six array sites. Each array site contains eight flash drives, and the HPFE Gen3 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.

4.5.7 RAID 6 overview

The DS8A00 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 while consuming less raw drive capacity than RAID 10.

Note: RAID 6 is the default and preferred array configuration in DS8A00 systems.

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 while 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 4-7 on page 131.

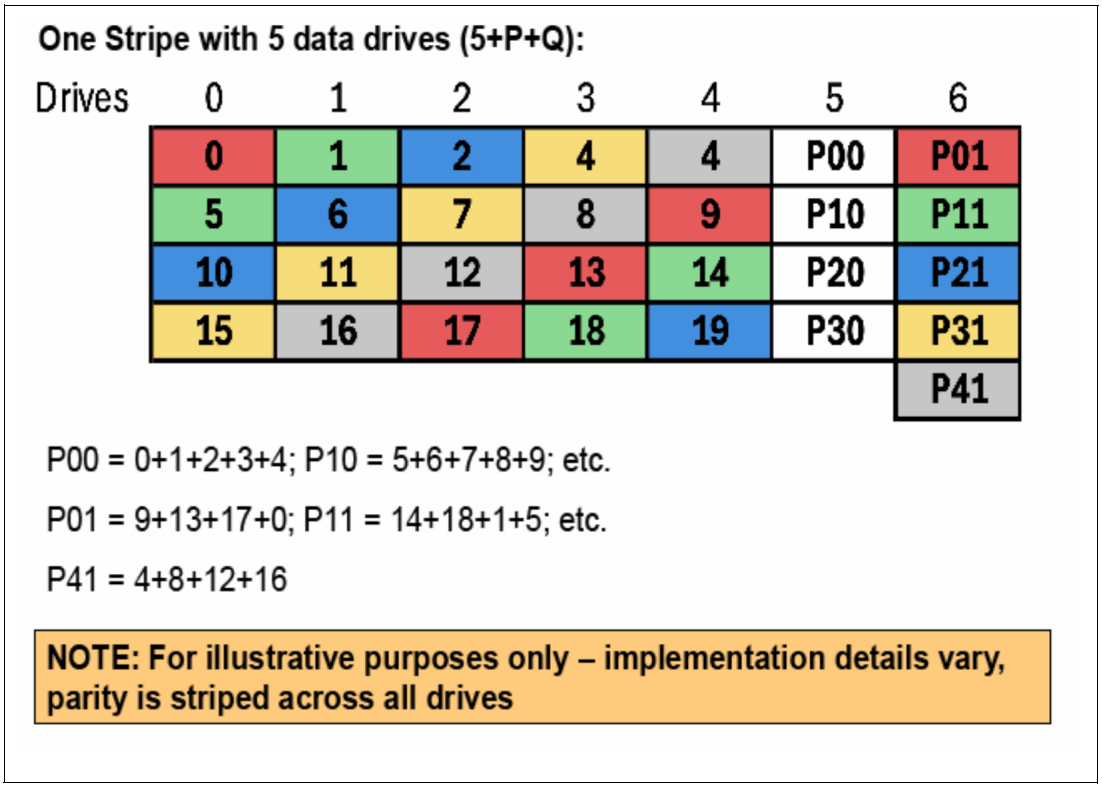


Figure 4-7 Illustration of RAID 6 stripe on a 5+P+Q+S array

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.

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 4-7.

RAID 6 implementation in the DS8A00

A RAID 6 array in one array site of a DS8A00 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 DS8A00, 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.

4.5.8 RAID 10 overview

RAID 10 provides high availability (HA) by combining the 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 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 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 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 DS8A00

In the DS8A00, 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 130.

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

4.5.9 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 other drives in this array encounter 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 NVMe 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 DS8A00 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 DS8A00 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 might, sometimes, be taken by a smaller drive array. Smart Rebuild corrects this situation after the failing drives are replaced, and returns the larger drive to the spare pool.

The Smart rebuild process avoids rebuilding 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.

4.5.10 Spare creation

This section describes methods of spare creation.

Flash drive enclosures

When the arrays are created on a DS8A00 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 DS8A00 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, DS8A00 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 4.5.5, “Disk scrubbing” on page 129 and 4.5.9, “Smart Rebuild” on page 133. This process does not fail the disk that is being moved. However, the process reduces the number of available spares in the DS8A00 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 flash drive is replaced with a new drive, the DS8A00 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 DS8A00 system because the drives are fully hot-pluggable. Each drive plugs into a NVMe 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.

4.5.11 Services Offering with Ransomware Threat Detection

FCMs provide real-time telemetry and analytics at the storage layer, monitoring I/O patterns and data entropy to detect ransomware-like behavior, such as unusual write spikes or access anomalies. By enabling Ransomware Threat Detection service, clients will receive ongoing updates throughout the subscription period and real-time notifications if a threat is detected.

Note: In Release 10.1, the notification will only be to IBM, the team training the AI model may reach out to the customer for more information on workload to tune the AI. Notification to the customer from DS8A00 will be enabled in the next release 10.2 with the acquired subscription.

For more information about Ransomware Threat Detection, see Chapter 14, “Ransomware Threat Detection” on page 487.

4.6 RAS on the power subsystem

This section describes the power subsystem components of the DS8A00 from a RAS standpoint.

All power and cooling components that constitute the DS8A00 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 2*N* 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 4.2.4, “NVS and power outages” on page 119).

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

4.6.1 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, HMC, and I/O enclosure power supply units

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, HMCs, CPCs, I/O enclosures, and Ethernet switches.

iPDUs are installed in pairs, one in each input power domain. An iPDU module can be replaced concurrently, as described in 2.6.2, “Power domains” on page 73.

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

For more information, see 2.6.1, “Intelligent Power Distribution Units” on page 69.

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 Storage DS8000 (10th Generation) Introduction and Planning Guide G10-I53-01.

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 Backup Power Modules (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 2.6.3, “Backup Power Modules and NVDIMM” on page 73.

The DS8A00 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 starts 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, HMC, and I/O enclosure power supply units

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

4.6.2 Line power loss

The DS8A00 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 DS8A00 uses NVDIMMs to preserve that data. For a full explanation of the NVS and cache operation, see 4.2, “CPC failover and failback” on page 116.

4.6.3 Line power fluctuation

If a power fluctuation occurs that causes a momentary interruption to power (often called a *brownout*), the DS8A00 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.

4.6.4 Power control

Power control functions are performed by the HMCs, which communicate sequencing information to the service processor in each CPC and iPDU. Power control of the DS8A00 can be performed by using the Service Maintenance Console Web User Interface (WUI) or by using the DS8000 Storage Management GUI or DS CLI commands. Figure 4-8 shows the Power Control in the Actions menu on the dashboard of the Storage Management GUI.

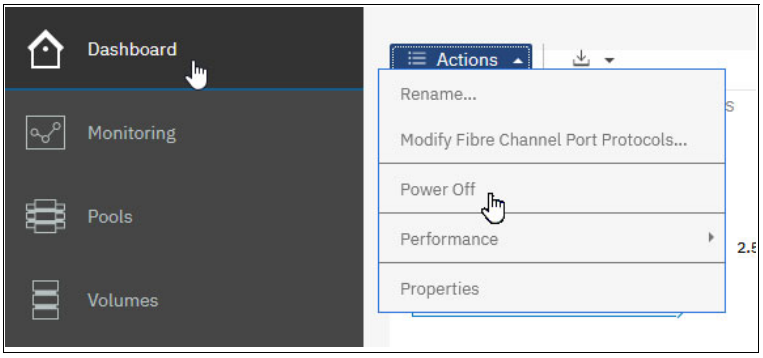


Figure 4-8 DS8A00 power control in the Storage Management GUI

Figure 4-9 shows the power control settings window of the Storage Management GUI.

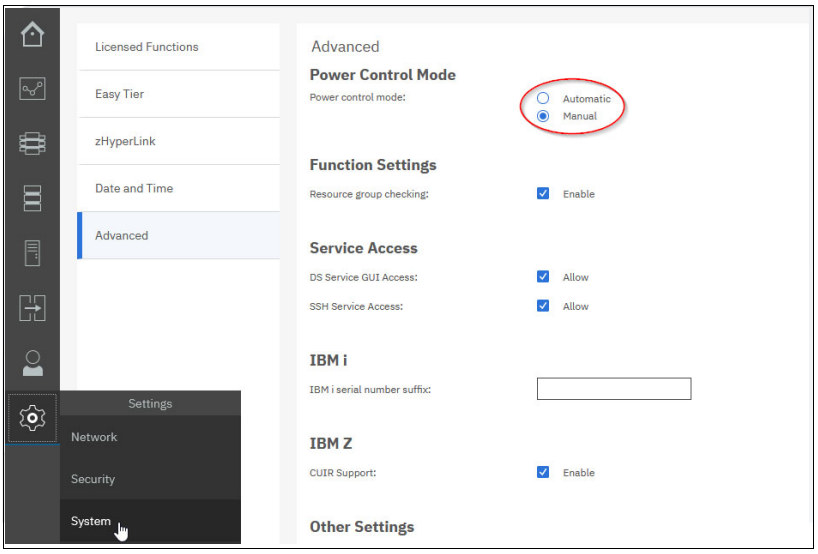


Figure 4-9 Modify power control settings from the Storage Management GUI

4.7 Other features

Many more features of the DS8A00 enhance RAS. Several of these features are described in this section.

4.7.1 Internal network

Each DS8A00 base frame contains two Gigabit Ethernet switches supporting a fully redundant pair of private management networks. Each CPC in the DS8A00 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 alternate private network. Section 6.1.2, “Private and Management Ethernet networks” on page 169 explains the design of the internal network in more detail.

Note: The Ethernet switches that are used internally in DS8A00 are for private network communication only. No external connection to the private networks is allowed. Client connectivity to the DS8A00 is allowed only through the provided external customer HMC Ethernet connectors (eth2 and eth1) at the rear of the base frame.

4.7.2 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 DS8A00. The kit is not available for the rack-mountable DS8A10 model A00. For more information, contact your IBM SSR.

4.7.3 IBM Certified Secure Data Overwrite

IBM Certified Secure Data Overwrite (SDO) is a process that provides a secure erasure of all data in a DS8A00 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 is completed. 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 4-10, the data becomes unintelligible after this process.

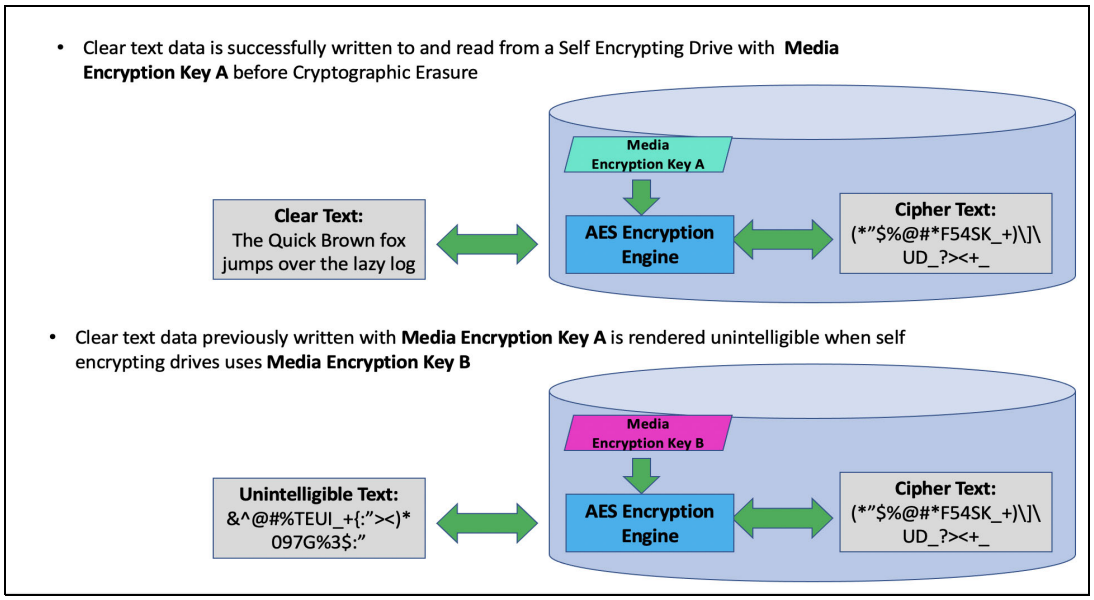


Figure 4-10 Secure Data Erasure data workflow

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. The HMC SSD 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-800-88. 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 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 drives are complete, the primary HMC drive 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 4-3 shows a sample SDO certificate.

Example 4-3 SDO sample certificate

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



Secure Data Overwrite Service for the IBM System Storage DS8A00

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 #: 5341-A05*75NKxxx
System Serial #: 75NKxxx


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


Date of Service Completion: Sep 25, 2025

In all cases, the successful complexion 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.

NVDIMMÕs 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
DS8A00 storage system with the serial number 75NHxxx



Drive Types Table
==========================
-------------------------------------------------------------------------------------
Drive Type       | Drive Block | Drive        | Number of    | Number of
2.5" FLASH-FDE   | Type        | Capacity     | drives       | drives available for
                 | 528         | (GB)         | installed    | 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 |ZAJ15KM0       |5000C5003023CB00     |R1-F07-D5       |Successful        |0                |
|2.5" FLASH-FDE |ZAJ15HT5       |5000C5003022849C     |R1-F07-D2       |Successful        |0                |
|2.5" FLASH-FDE |ZAJ15GWJ       |5000C50030228068     |R1-F07-D1       |Successful        |0                |
|2.5" FLASH-FDE |ZAJ15L1R       |5000C5003023CBB8     |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 DS8A00 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 | hdisk  | CPC Drive     |  Overwrite     |  Completion          |
| Complex # | Number | Serial Number |  Status        |  Date                |
-------------------------------------------------------------------------------
|  CPC 0    | hdisk0 | WAE1045Q      |  Successful    |  2025/09/24 19:11:46 |
|  CPC 0    | hdisk1 | WAE10N39      |  Successful    |  2025/09/24 20:10:31 |
|  CPC 1    | hdisk0 | 0TJ5SJLP      |  Successful    |  2025/09/24 19:13:16 |
|  CPC 1    | hdisk1 | WAE104DZ      |  Successful    |  2025/09/24 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  |            N/A             | Successful  | 25/09/25-06:33:51 | SSD              |
|     Console        |                            |             |                   |                  |
-------------------------------------------------------------------------------------------------
|Secondary Management|            N/A             | Successful  | 25/09/25-03:49:39 | SSD              |
|     Console        |                            |             |                   |                  |
-------------------------------------------------------------------------------------------------

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   | Overwrite  | Completion           |
|                          |               | Capacity  | Status     | Date                 |
===============================================================================================
| U78D2.001.WZS03ED-P1-C22 | YH10DP98G0E7  |  32 GB    | Successful | 25/09/24 19:42:56 |
| U78D2.001.WZS03ED-P1-C36 | YH10DP98G0EX  |  32 GB    | Successful | 25/09/24 19:42:56 |
===============================================================================================
===============================================================================================
| U78D2.001.WZS03RG-P1-C22 | YH30DP944055  |  32 GB    | Successful | 25/09/24 19:42:58 |
| U78D2.001.WZS03RG-P1-C36 | YH30DP9440OR  |  32 GB    | Successful | 25/09/24 19:42:58 |
===============================================================================================
```

5

IBM DS8000 physical planning and installation

This chapter describes the various steps that are involved in the planning and installation of the IBM DS8A00. 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” on page 144
- “Planning for the physical installation” on page 146
- “Network connectivity planning” on page 159
- “Remote Mirror and Remote Copy connectivity” on page 163
- “Disk capacity considerations” on page 163
- “z/OS maintenance” on page 165

For more information about the configuration and installation process, see IBM Storage DS8000 (10th Generation) Introduction and Planning Guide, G10-I53-01.

© Copyright IBM Corp. 2025. All rights reserved.

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: Both the IBM DS8A50 and IBM DS8A80 systems can support an expansion frame that can be installed adjacent (Feature Code AB00) or 20 meters away from the base frame (Feature Code AB01).

- 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 DS8000. If encryption activation is required, consider either local or external key management. If external, then consider the location and connection needs for the external key servers, such as IBM Guardium 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 6.6.2, “Remote authentication” on page 186.
- 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, “DS8000 G10 configuration flow” on page 221.

5.1.1 Client responsibilities for the installation

The DS8000 is specified as an IBM or IBM Business Partner installation and setup system. However, the client is responsible for several high-level planning and installation activities, including the following:

- Physical configuration planning

  Your Storage Marketing Specialist can assist in the planning and selecting the physical configuration and features of the DS8A00 model.

- Installation planning
- Integration of LDAP

  IBM can assist in the planning and implementation at the client’s request.

- Connectivity for Assist On-site (AOS)

  IBM can assist in the planning and implementation 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 Guardium Key Lifecycle Manager, Gemalto Safenet KeySecure, and Thales Vormetric DSM. IBM 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 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 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 are also client responsibilities. The logical configuration can be created, applied, and modified by using the DS GUI, or DS Command-line Interface (DS CLI).

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 DS8000 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 DS8000 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 DS8000 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 DS8A00 models

Shipping container	Packaged dimensions (in meters and inches)	Maximum packaged weight (in kilograms and pounds)
IBM DS8A10 model A00	3 Pallets each of: Height 1.32 m (51.96 in.) Width 1.15 m (45.27 in.) Depth 0.14 m (5.51 in.)	825.8 kg (1820.57 lb)
IBM DS8A10 model A01	Height 2.22 m (87.7 in.) Width 1.00 m (39.4 in.) Depth 1.50 m (59.1 in.)	924 kg (2037 lb)
DS8A50 model A05	Height 2.22 m (87.7 in.) Width 1.00 m (39.4 in.) Depth 1.50 m (59.1 in.)	924 kg (2037 lb)
DS8A80 model A08	Height 2.22 m (87.7 in.) Width 1.00 m (39.4 in.) Depth 1.50 m (59.1 in.)	924 kg (2037 lb)
Expansion Frame model E05	Height 2.22 m (87.7 in.) Width 1.00 m (39.4 in.) Depth 1.50 m (59.1 in.)	603 kg (1330 lb)

5.2.2 Floor type and loading

The DS8000 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 DS8A00 models, see Table 5-1.

Important: Verify 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 DS8000. 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.)

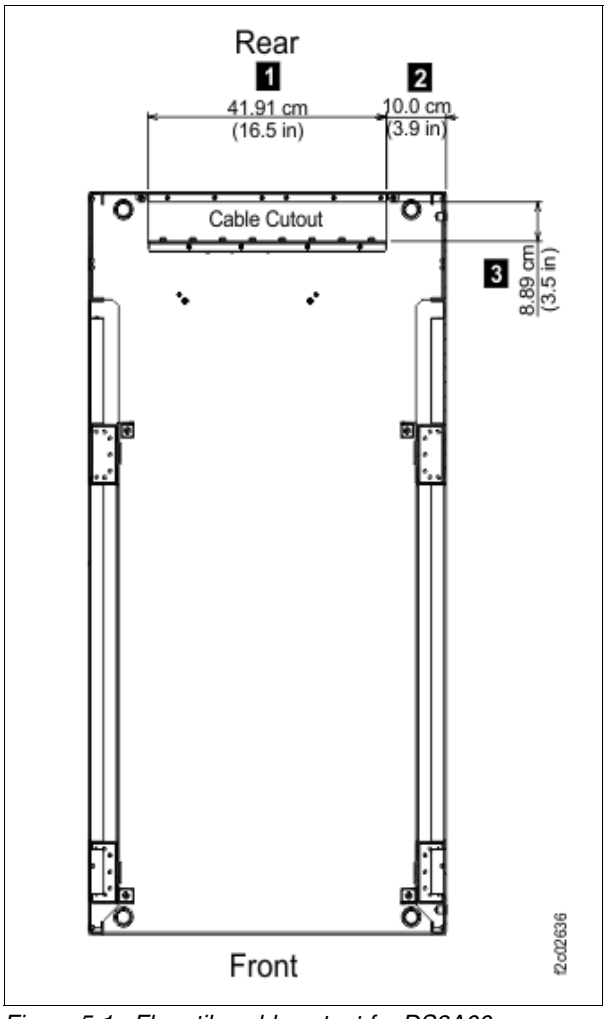


Figure 5-1 Floor tile cable cutout for DS8A00

For more information about floor loading and weight distribution, see *IBM Storage DS8000 (10th Generation) Introduction and Planning Guide,* G10-I53-01.

5.2.3 Overhead cabling features

The overhead cabling (top exit) feature, as shown in Figure 5-2 on page 148, is available for the DS8000 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 AAV1 Top exit bracket for overhead cabling
- Feature Code AAX0 Overhead ladder

For more information, see *IBM Storage DS8000 (10th Generation) Introduction and Planning Guide,* G10-I53-01.

Figure 5-2 shows the overhead cabling (top exit) feature.

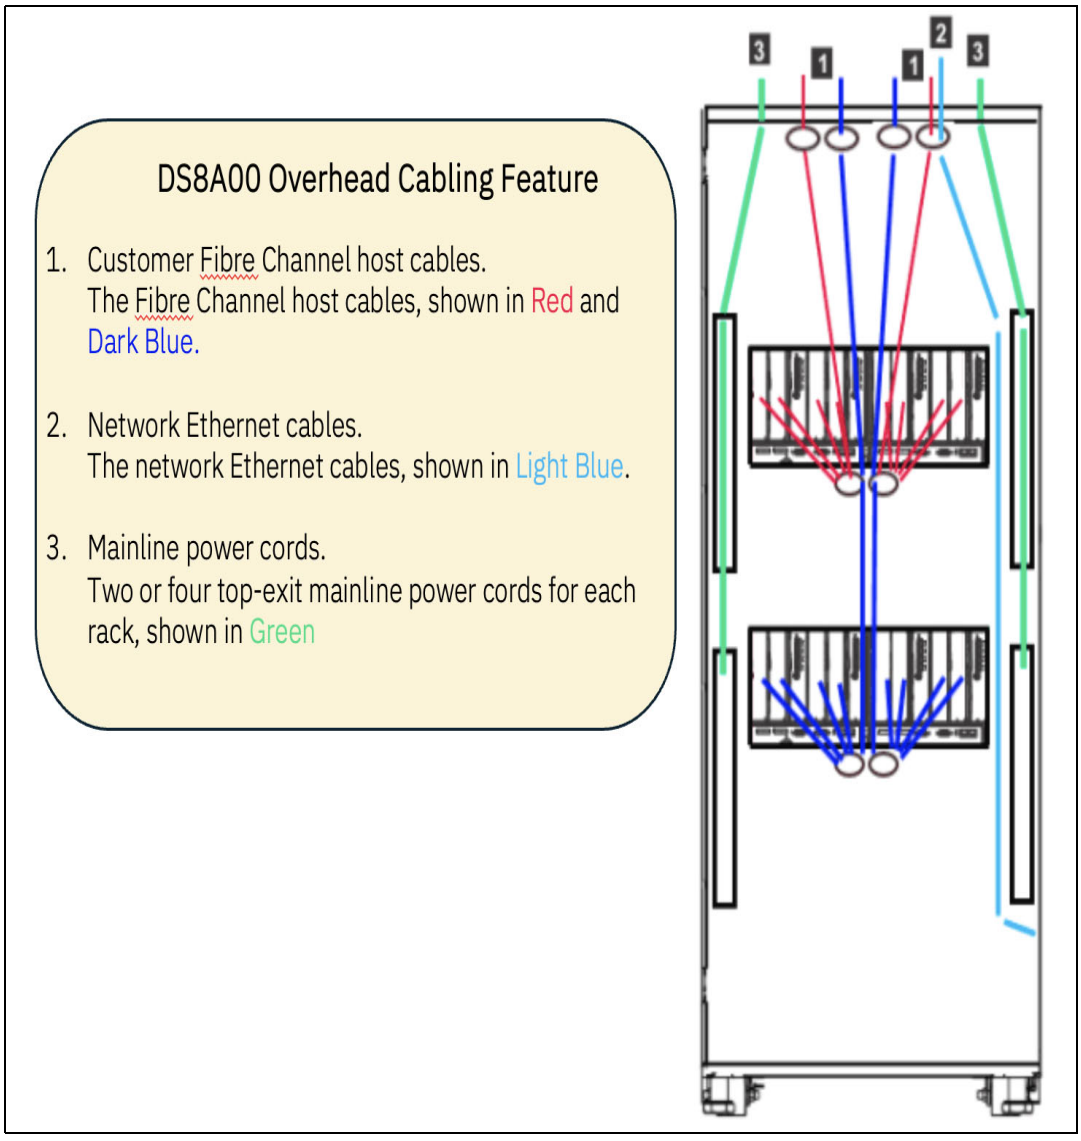


Figure 5-2 Overhead cabling for DS8A00

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 DS8A00 dimensions with decorative doors

Dimensions with casters and covers	DS8A00 all models (racked)
Height	204 cm (80.3 in.)
Width	62 cm (24.4 in.)
Depth	124 cm (48.8 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 DS8A00 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.

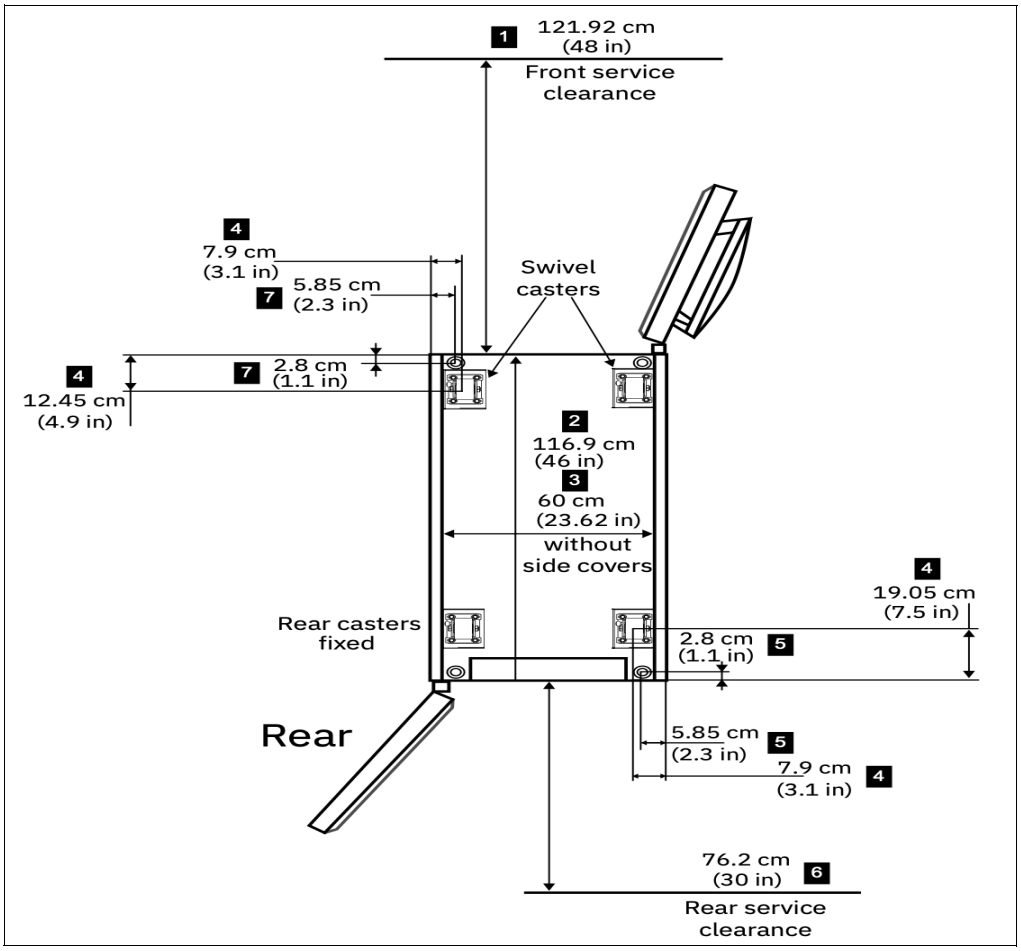


Figure 5-3 DS8A00 service clearance requirements

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 DS8000 power requirements:

- Power connectors
- Input voltage
- Power consumption and environment
- Power control features

Power connectors

Each DS8000 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 has 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 Storage DS8000 (10th Generation) Introduction and Planning Guide,* G10-I53-01.

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 DS8A00 input voltages and frequencies.

Table 5-3 DS8A00 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	30–32 Amps[a]
Steady-state input frequency	50 ± 3 or 60 ± 3 Hz
Power line disturbance (PLD) input frequencies (<10 seconds)	50 ± 3 or 60 ± 3 Hz

a. Can vary by region.

Table 5-4 DS8A00 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	30–32 Amps
Steady-state input frequency	50 ± 3 or 60 ± 3 Hz	
PLD input frequencies (<10 seconds)	50 ± 3 or 60 ± 3 Hz	

Power consumption

Table 5-5 lists the power consumption specifications of the DS8A00. 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)	DS8A10 model A00: 7.9 DS8A10 model A01: 7.9 DS8A50 model A05: 9.1 DS8A10 model A08: 9.1	Model E05: ≤10.2
Thermal load	British thermal units (BTUs) per hour	≤34800	≤ 34800
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 model A01, A05, and E01 that contained four fully configured High Performance Flash Enclosure Gen3 Pairs and 16 Fibre Channel adapters.
- All Frames and configurations were used in single-phase mode and did not exceed 8 kVA.

DS8A00 cooling

Air circulation for the DS8000 is provided by the various fans that are installed throughout the frame. All of the fans in the DS8000 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.

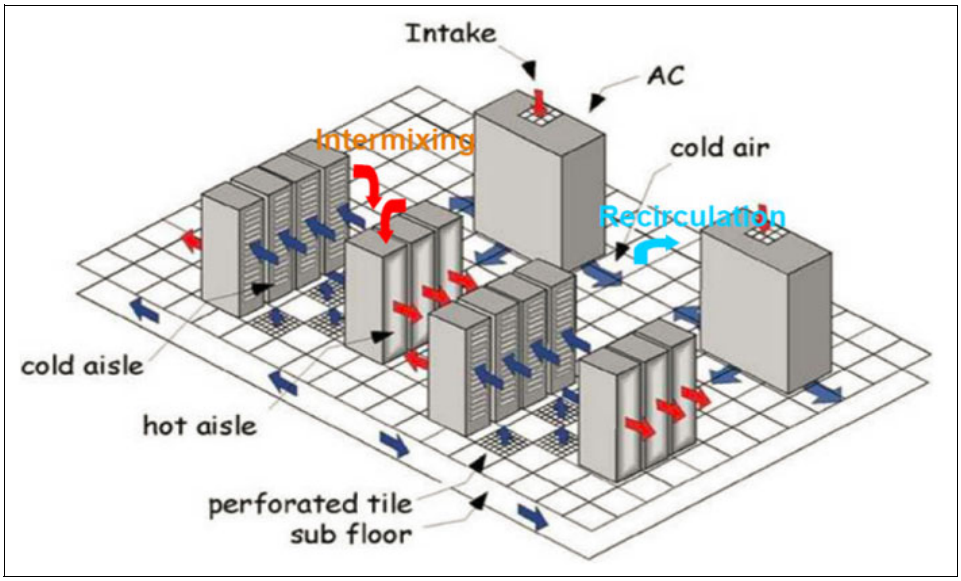


Figure 5-4 DS8A00 air flow: Hot aisle/cold aisle approach

The operating temperature for the DS8A00 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 DS8000 system is installed:

- Ensure that the air circulation for the DS8000 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 DS8000 system.

Power control features

The DS8A00 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 Storage DS8000 (10th Generation) Introduction and Planning Guide,* G10-I53-01.

5.2.6 Host interface and cables

Both DS8A10 models, DS8A50 model A05, and DS8A80 Model A08 can contain a maximum of 16 host adapters. The DS8A50 and DS8A80 model E05 supports an extra 16 host adapters to be installed in the expansion frame for a total of 32 host adapters. For a full breakdown of the available ports on a DS8A00, see Table 5-6.

Table 5-6 shows the minimum and maximum numbers of host adapters that are supported by the DS8A00.

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
DS8A10 model A00 and DS8A10 model A01	Base	2	16
DS8A50 model A05 and DS8A80 model A08	Base frame	2	16
DS8A50 and DS8A80 model E05	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 µm 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 maximum supported distance with Fibre Channel speeds

Fibre Channel Speed	OM3	OM4	OM5
16Gbps	100m	125m	150m

Fibre Channel Speed	OM3	OM4	OM5
32Gbps	70m	100m	150m
64Gbps	35m	50m	100m

The FC and FICON longwave (LW) host adapter, when it is used with 9 µm 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
AB10	40 m (131 ft.)	50 µm multimode	► SW FC or FICON host adapters (Feature Code AAE0) ► LC connector
AB13	3 m (10 ft.)	50 µm multimode	
AB11	31 m (102 ft.)	50 µm multimode	► LC/SC connector
AB20	31 m (102 ft.)	9 µm single mode	► LW FC or FICON host adapters (Feature Code AAE1) ► LC connector
AB23	3 m (10 ft.)	9 µm single mode	
AB21	31 m (102 ft.)	9 µm single mode	► LC/SC connector

For more information about host types, models, adapters, and operating systems (OSs) that are supported by the DS8A00, 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) Gen4 with a maximum distance of 150 meters. zHyperLink connects the IBM Z (central processor complexes (CPCs)) directly to the I/O enclosure of a DS8A00 system, as shown in Figure 5-5 on page 154.

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 DS8A00, 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.

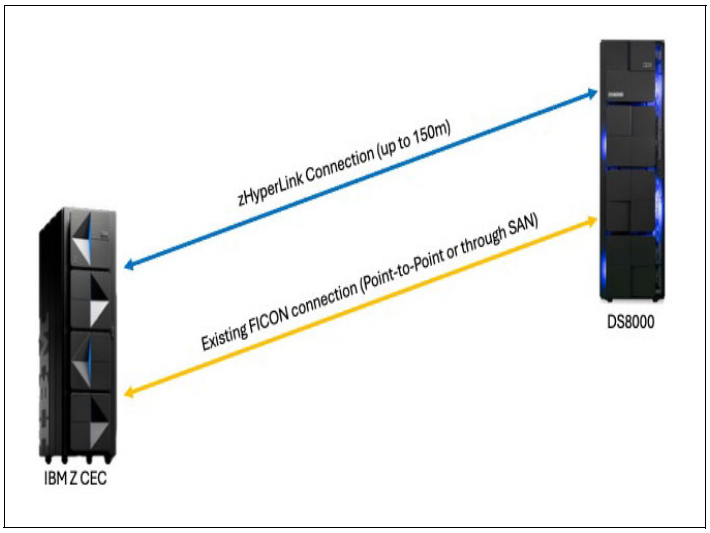


Figure 5-5 zHyperLink connection

The number of zHyperLinks that can be installed based on the number of cores available is listed in Table 5-9.

Table 5-9 zHyperLink availability by DS8A00 model

System or model	Cores per CPC (DS8A00 server)	zHyperLink support	Max zHyperLink connections (increments of 2)
DS8A10 model A00	10	Yes	4
DS8A10 model A01	10	Yes	4
DS8A50 model A05	40	Yes	8
DS8A80 model A08	48	Yes	8
DS8A50 model E05	-	Yes	10
DS8A80 model E05	-	Yes	12

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
AAG0	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
AB30	OM4 50/125 micrometer, multimode, and MTP connectors	40 m (131 Ft)	zHyperLink I/O adapter (Feature Code AAG0)
AB31		150 m (492 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 DS8000 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 DS8000 provides the following configuration capabilities:

- Fabric or point-to-point topologies
- A maximum of number of host adapter ports that depends on the model:
  - 64 ports on both DS8A10 Models, DS8A50 model A05, and DS8A80 model A08
  - 64 ports on the DS8A50 and DS8A80 model E05 expansion frame, 128 ports with the B 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 z17 and z16® servers support 32,000 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 DS8A00 supports 32 GFC 4-port host adapters.

> **Note:** IBM z17 supports 32 GFC host adapters and FICON Express32-4P and Express32S channels. The support for 32 GFC host adapters provides reduced latency for large read/write operations and increased port bandwidth compared to 16 GFC host adapters on previous models, thus taking full advantage of 32 GFC host adapters on the DS8A00.

- 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 .

> **Note:** DS8000 has a set of internal parameters that are known as *pokeables*, which are sometimes 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 DS8000 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 DS8000 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

```
dscli> lssi
Date/Time: October 22, 2025 at 10:20:16 AM CEST IBM DSCLI Version: 7.10.10.348 DS: -
Name        ID               Storage Unit     Model WWNN             State  ESSNet
==========================================================================================
DS8A10-DS02 IBM.2107-78xxxx1 IBM.2107-78xxxx0 A01   5005076303FFF1D5 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

```
dscli> lssu
Date/Time: October 22, 2025 at 10:22:49 AM CEST IBM DSCLI Version: 7.10.10.348 DS: -
Name        ID               Model WWNN            pw state
=============================================================
DS8A10-DS02 IBM.2107-78xxxx0 A01   5005076303FFC9D5
On
```

Determining a WWPN by using a DS CLI

Similar to the WWNN, a WWPN in the DS8A00 looks like the following address:

`50:05:07:63:0z:YY:Yx:xx`

However, the DS8000 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 **lsioport** to determine the SFI WWPN, as shown in Example 5-3.

Example 5-3 WWPN determination

```
dscli> lsioport
Date/Time: October 22, 2025 at 10:24:49 AM CEST IBM DSCLI Version: 7.10.10.348 DS:
IBM.2107-78xxxx1
ID    WWPN             State                      Type             topo     portgrp Security
==========================================================================================
I0020 50050763030231D5 Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
I0021 50050763030271D5 No light detected         Fibre Channel-SW SCSI-FCP 0       Disabled
I0022 500507630302B1D5 Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
I0023 500507630302F1D5 No light detected         Fibre Channel-SW SCSI-FCP 0       Disabled
I0120 50050763030A31D5 Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
I0121 50050763030A71D5 No light detected         Fibre Channel-SW SCSI-FCP 0       Disabled
I0122 50050763030AB1D5 Communication established Fibre Channel-SW FICON    0       Disabled
I0123 50050763030AF1D5 No light detected         Fibre Channel-SW SCSI-FCP 0       Disabled
I0220 50050763031231D5 Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
I0221 50050763031271D5 Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
I0222 500507630312B1D5 No light detected         Fibre Channel-SW SCSI-FCP 0       Disabled
I0223 500507630312F1D5 Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
I0230 50050763031331D5 Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
I0231 50050763031371D5 No light detected         Fibre Channel-SW SCSI-FCP 0       Disabled
I0232 500507630313B1D5 Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
I0233 500507630313F1D5 No light detected         Fibre Channel-SW SCSI-FCP 0       Disabled
I0320 50050763031A31D5 Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
I0321 50050763031A71D5 No light detected         Fibre Channel-SW SCSI-FCP 0       Disabled
I0322 50050763031AB1D5 Communication established Fibre Channel-SW FICON    0       Disabled
I0323 50050763031AF1D5 Communication established Fibre Channel-SW SCSI-FCP 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.

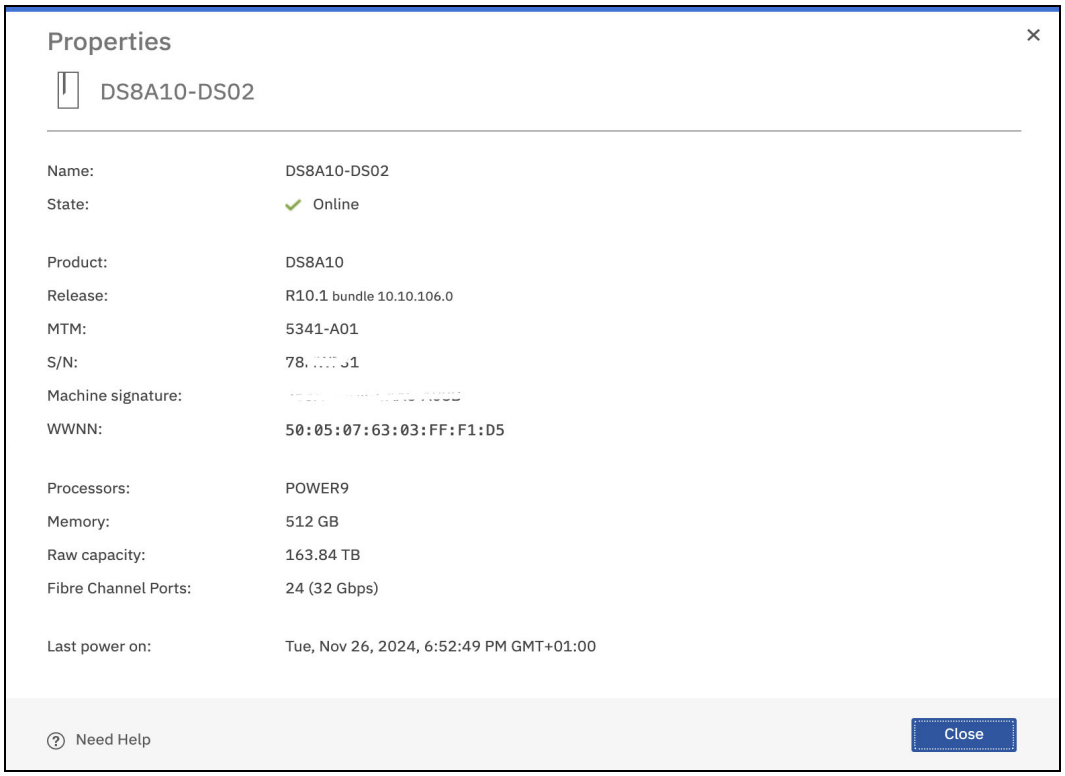


Figure 5-6 System properties: WWNN

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.

4. 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 on page 159.

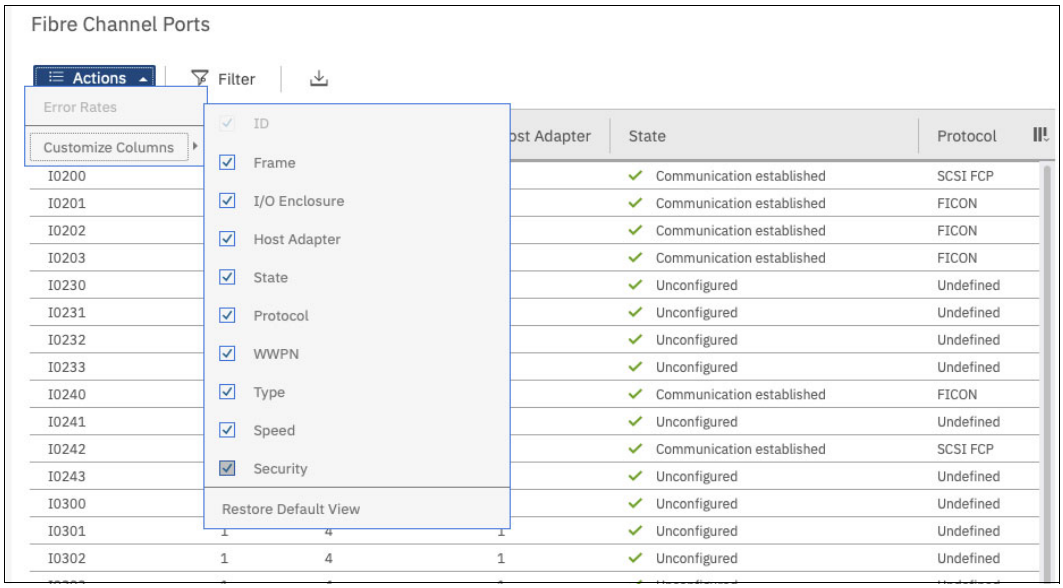


Figure 5-7 Determining the I/O port WWPN

5.3 Network connectivity planning

To implement a DS8000, 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.
- A IBM 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 Storage DS8000 (10th Generation) Introduction and Planning Guide,* G10-I53-01.

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 DS8000 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 can also control a second Ethernet adapter within the HMCs.

For more information about HMC planning, see Chapter 6, “IBM DS8A00 Management Console planning and setup” on page 167.

Important: The DS8000 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 12.12, “Using IBM Storage Insights” on page 451.

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 DS8000 over the client's network. For more information about hardware and software requirements for the DS CLI, see *IBM DS8000 (10th Generation) Series Command-Line Interface User's Guide*, G10-C53-01.

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. The DS8000 uses an embedded AOS server solution. Embedded AOS is a secure and fast broadband form of remote access.

IBM Remote Support Center

The DS8000 also supports a web-based remote support option that is called RSC. It provides more isolation between IBM Support and the DS8000 system by requiring that all remote support actions be performed by using a web-based console interface.

For more information, see Chapter 6, "IBM DS8A00 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 DS8000 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 DS8A00 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 a DS8A00 system, you can choose among IBM 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 Guardium Key Lifecycle Manager. You cannot mix IBM 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 10.0)*, REDP-4500-11.

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 Guardium Key Lifecycle Manager

An IBM Guardium Key Lifecycle Manager license is required to use the IBM Guardium Key Lifecycle Manager software. Two isolated Guardium Key Lifecycle Manager servers are required to enable data at rest encryption on the DS8A00 system.

Important: Clients must acquire an IBM Guardium Key Lifecycle Manager license to use the Guardium Key Lifecycle Manager software.

The licensing for IBM 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 is supported with current IBM Guardium Key Lifecycle Manager V5.x versions. 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.

You are advised to upgrade to the latest version of the IBM Guardium Key Lifecycle Manager.

IBM Guardium Key Lifecycle Manager connectivity and routing information

To connect the IBM Guardium Key Lifecycle Manager to your network, provide the following settings to your IBM SSR:

- IBM 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 DS8000 connection and to obtain an administration management interface to the IBM Guardium Key Lifecycle Manager server. These ports are defined by the IBM Guardium Key Lifecycle Manager administrator.

For more information, see the following IBM publications for IBM 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 Guardium Key Lifecycle Manager 5.x

Local Key Manager

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 10.0)*, REDP-4500-11.

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.

LDAP authentication can be configured through the Storage Management GUI, as described in 6.6.2, “Remote authentication” on page 186.

5.4 Remote Mirror and Remote Copy connectivity

The DS8000 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 10.1*, SG24-8367-02.

5.5 Disk capacity considerations

The effective capacity of a DS8A00 is determined by the following factors:

- The spare configuration
- The capacity of the installed drives
- The selected RAID configuration: RAID 6 or RAID 10
- 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 Gen3

High-Performance Flash Enclosures (HPFEs) are installed in pairs, with 16, 32, or 48 NVMe flash drives per enclosure pair. Two spare NVMe flash drives are assigned for each HPFE Gen3 pair. If a flash drive fails and a spare is taken, the system calls for service because only one spare remains in the HPFE Gen3 pair (DA pair).

For more information about the DS8000 sparing concepts, see 4.5.10, “Spare creation” on page 134.

5.5.2 Disk capacity

The following RAID configurations are supported on the DS8A00:

- 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.

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 DS8000 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 Storage DS8000 (10th Generation) Introduction and Planning Guide,* G10-I53-01

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 set of 16 disk drives.

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.

The DS8A00 offers the following flash drives sets with HPFE Gen3:

- 2.5-inch high-capacity flash (Tier 1) drives are 1.6 TB.
- 2.5-inch high-performance flash (Tier 2) drives are 1.92 TB, 3.84 TB, and 7.68 TB capacity drives.

Flash drives in HPFE Gen3 are ordered in sets of 16 within an enclosure pair. There are three sets of 16 drives in an HPFE Gen3 enclosure pair.

Important: The following restrictions apply:

- Within one HPFE Gen3 pair of six array sites, a RAID intermix is allowed, but no intermix of Capacity tier drives with Compression tier drives (Capacity tier 1 or Capacity tier 2) is supported.
- For the latest information about supported RAID configurations and requesting an RPQ or SCORE, contact your IBM SSR.

5.6 z/OS maintenance

The `MISSINGFIX` report is required to identify if any APARs are needed to use a function that is not yet installed. The `REPORT MISSINGFIX` command checks the specified zones (using the zones operand) and determines if any fixes are missing based on the specified fix categories.

Note: The DFSMSdss ((Data Facility Storage Management Subsystem Data Set Services) Compression tool utility is included in z/OS 3.2 and z/OS 3.1 but not included in z/OS 2.5 release. Therefore, all applicable APARs and PTFs must be installed.

Run the SMP/E **`REPORT MISSINGFIX`** for the following function fix categories:

- **`IBM.Function.DiskCompression`**

  Enables or exploits hardware compression on the IBM DS8A00 series devices.
- **`IBM.Function.zHyperLink`**

  Fixes for the zHyperLink function.
- **`IBM.Function.DFSMSCloudStorage`**

  Fixes for Transparent Cloud Tiering Function.

Example 5-4 shows a sample **`MISSINGFIX`** JCL.

Example 5-4 Sample SMP/E Control cards with FIXCAT option

```
SET BOUNDARY (GLOBAL) .
REPORT
MISSINGFIX
ZONES (<your zone names>)
FIXCAT(IBM.Function.DiskCompression,
IBM.Function.zHyperLink,
IBM.Function.DFSMSCloudStorage) .
```

For more information about the SMP/E APPLY or REPORT MISSINGFIX commands, See SMP/E for z/OS commands.

6

IBM DS8A00 Management Console planning and setup

This chapter describes the planning tasks that are involved in the setup of the IBM DS8A00 Management Console (MC), which is usually referred to as the Hardware Management Console (HMC).

This chapter covers the following topics:

- “DS8A00 Management Console overview” on page 168
- “Management Console software” on page 172
- “Management Console activities” on page 179
- “Management Console network settings” on page 182
- “User management” on page 184
- “Secondary Management Console” on page 194

© Copyright IBM Corp. 2025. All rights reserved.

6.1 DS8A00 Management Console overview

The HMC 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 an interface where service personnel perform diagnostic and repair tasks.

The HMC 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 HMC is a configuration and management station for the whole DS8A00 system.

To enhance security by ensuring that only trusted software is run, starting with DS8000 release 10, the HMC uses Secure Boot.

The DS8A00 includes two HMCs, for redundancy reasons.

The HMC, which is the focal point for DS8A00 management, includes the following functions:

- DS8A00 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
- Remote support access capability
- 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 Console hardware

The HMC itself consists of a 7063-CR2 Power HMC running Red Hat Linux. It supports the DS8A00 hardware and firmware installation and maintenance activities.

Both of the 1U primary and secondary HMCs are mounted in the bottom of the main rack. The primary HMC is installed above the secondary HMC. There is an 1U keyboard and display tray that are available.

The HMC connects to the customer network and provides access to functions that can be used to manage the DS8A00. 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 DS8A00.

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 HMC 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 HMCs are equipped with Ethernet connections for the client's network. For more information, see 6.1.2, "Private and Management Ethernet networks" on page 169.

To provide continuous availability to the HMC functions, the DS8A00 includes a second HMC. The secondary HMC is needed for redundancy, such as for encryption management or CS management functions. For more information about the secondary HMC, see 6.7, "Secondary Management Console" on page 194.

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

```
dscli> lshmc
Date/Time: September 29, 2025 at 1:39:34 PM EDT IBM DSCLI Version: 7.10.10.348 DS: -
Name          State  Role         Release                  Management-IP Location-Code    Machine-Type Machine-Model Memory Disk
=============================================================================================================================
DS8A50-1A.DMZ Online Primary(1)   R10.1 bundle 10.10.106.0 192.168.61.90 U7063CR2.78A81DD 7063         CR2           33GB   1TB
DS8A50-1B.DMZ Online Secondary(2) R10.1 bundle 10.10.106.0 192.168.61.91 U7063CR2.78A97DD 7063         CR2           33GB   1TB
dscli>
```

6.1.2 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 Rack Control Hubs (RCH) that are isolated from each other at the bottom of the rack.

Each central processor complex (CPC) flexible service processor (FSP) and each CPC logical partition (LPAR) network are connected to both RCHs. Each of these components (FSP and LPAR) uses their own designated interface for the black network and another interface for the gray network.

Each HMC also uses two designated Ethernet interfaces for the internal black (eth0) and gray (eth3) networks.

Additionally, an HMC contains two Ethernet interfaces (eth1 and eth2) for the customer network connection to allow management functions to be started over the network. These adapters 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 6-1 on page 170 and Table 6-1 on page 170 show these internal and external network connections.

Important: The Rack Control Hubs that are shown in Figure 6-2 on page 171 (the black and gray private networks) are for DS8A00 internal communication only. Do not connect these ports directly to your network. There is no connection between the customer network interfaces and the isolated black and gray networks.

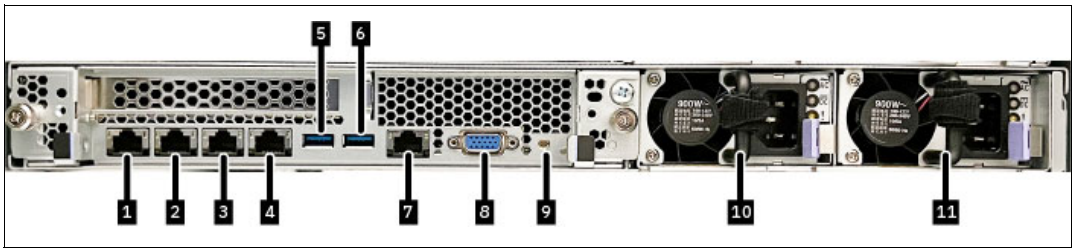


Figure 6-1 HMC network connections

Table 6-1 Network connections

From HMC network connections	Connects to
1	DS8000 private (gray) network (eth3)
2	Primary customer network connection (eth2)
3	Secondary customer network connection (optional) (eth1)
4	DS8000 private (black) network (eth0)
5	USB 2.0 serial communication to rack control hub - USB port 4
6	USB 2.0 to keyboard cable couplers (Note 1) - USB port 3
7	Direct BMC port (maintenance) - IPMI (not used in DS8K)
8	VGA to monitor cable couplers - VGA
9	HMC status LED
10	Power supply unit 1 (left, labeled E0)
11	Power supply unit 0 (right, labeled E1)

Intelligent Power Distribution Units

In DS8A00 systems, the rack power control is managed through the Ethernet-managed intelligent Power Distribution Units (iPDUs).

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 black and gray Rack Control Hubs. They are distributed over the black and gray networks, 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 RCH and the iPDUs that belong to the other power domain (usually on the right side of the rear of the rack) connect to the black network RCH.

There are two10-port Rack Control Hubs. One is for the black network, and one is for the gray network. The 1U space that is required is already reserved at the bottom of the base rack.

Figure 6-2 and Table 6-2 show how each port is used on the pair of DS8A00 Rack Control Hubs.

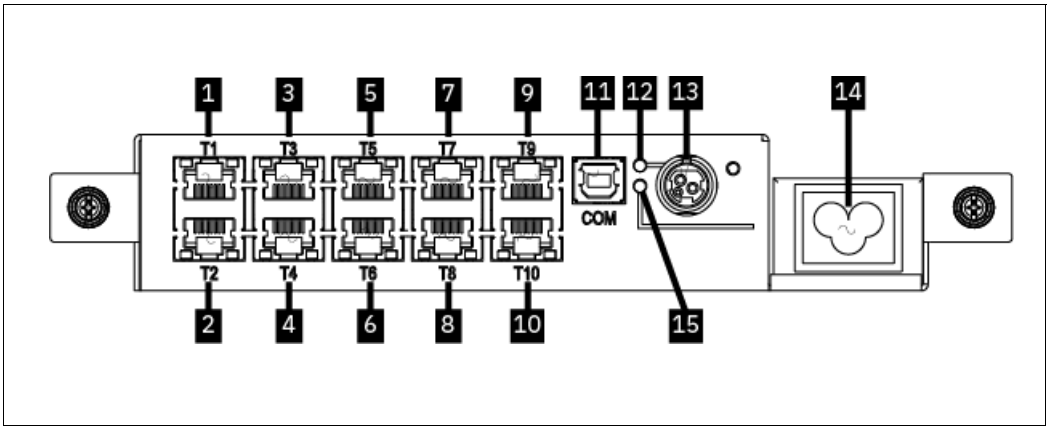


Figure 6-2 Rack Control Hub breakout

Table 6-2 Black network Rack Control Hub connections

Left side RCH (Gray)	Rack Control Hub Connection	Right side RCH (Black)
Upper FSP	1	Upper FSP
Lower FSP	2	Lower FSP
Upper LPAR	3	Upper LPAR
Lower LPAR	4	Lower LPAR
HMC1 (Upper)	5	HMC1 (Upper)
HMC2 (Upper)	6	HMC2 (Upper)
Rack 1 iPDU, upper left	7	Rack 1 iPDU, upper right
Rack 2 iPDU, lower left	8	Rack 2 iPDU, lower right
Rack 2 iPDU, upper left	9	Rack 2 iPDU, upper right
Unused	10	Unused
Serial communication (SWn-USB COM)	11	Serial communication (SWn-USB COM)
Partner rack control hub power LED	12	Partner rack control hub power LED
Partner power cable (SWn-PTNR PWR)	13	Partner power cable (SWn-PTNR PWR)
PSU / input power cord	14	PSU / input power cord
Local rack control hub power LED	15	Local rack control hub power LED

6.2 Management Console software

The HMC, 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, external Copy Services Manager, and the DS CLI to connect to the DS8A00 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 DS8A00 HMC
- 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 231.

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 HMC. 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 DS8A00 Storage Management Command-line Interface” on page 335. 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

DS8A00 RESTful API services provide an application programming interface (API) to manage DS8A00 through communication with the HMC. The RESTful API communicates with RESTful services that run on the HMC. The RESTful services in turn interact with the

IBM ESSNI server software that runs on the HMC to pass requests and receive replies. For more information about the RESTful API, see IBM DS8900F/DS8880 RESTful API Guide V4.0 SC27-9823-00.

6.2.4 IBM Copy Services Manager interface

IBM Copy Services Manager is preinstalled on the DS8A00 HMC. 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 DS8A00 code maintenance schedule, which is performed by IBM support personnel, or with the DS CLI as explained in 6.2.5, “Updating the embedded IBM Copy Services Manager” on page 173

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 the DS CLI, if needed.
5. Update IBM Copy Services Manager on the HMC.

The DS8000 Code Recommendation page provides a link to the DS8A00 code bundle information page, as shown in Figure 6-3 on page 174 and Figure 6-4 on page 174.

Model	Installed Level	"go to" Level	Latest Level
DS8A00 (R10.1)	R10.1 GA - 10.10.106.0 [Sep 2025] (62, **65**)	R10.1 GA - 10.10.106.0 [Sep 2025] (62, **65**)	R10.1 GA - 10.10.106.0 [Sep 2025] (62, **65**)
DS8A00 (R10.0)	R10.0 SP1 - 10.0.245.0 [Nov 2024] (61, 62, 63, 64)	R10.1 GA - 10.10.106.0 [Sep 2025] (62, **65**)	R10.0 SP2.2 - 10.2.45.0 [Aug 2025] (61, 62)

Figure 6-3 DS8A00 Code Recommendation page

Release 10.1

DS8A00 Bundle	SEA or LMC Version	DSCLI Client	Heat Map Transfer Utility	Storage Manager	Copy Services Manager	Threat Detection
10.10.106.0	7.10.10.348	7.10.10.348	7.10.10.348	5.10.10.1115	6.3.15	0.0.0.13

Release 10.0

DS8A00 Bundle	SEA or LMC Version	DSCLI Client	Heat Map Transfer Utility	Storage Manager	Copy Services Manager
10.2.45.0	7.10.2.131	7.10.2.131	7.10.2.131	5.10.2.1048	6.3.14
10.2.43.0	7.10.2.127	7.10.2.127	7.10.2.127	5.10.2.1046	6.3.14
10.2.35.0	7.10.2.111	7.10.2.113	7.10.2.113	5.10.2.1038	6.3.14
10.1.3.0	7.10.1.8	7.10.1.8	7.10.1.8	5.10.1.1003	6.3.12
10.0.245.0	7.10.0.882	7.10.0.882	7.10.0.882	5.10.0.1297	6.3.12
10.0.237.0	7.10.0.861	7.10.0.861	7.10.0.861	5.10.0.1289	6.3.12

Figure 6-4 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 **lssoftware** DS CLI command:

```
lssoftware -l -type csm -hmc all
```

Example 6-2 shows an example where the IBM Copy Services Manager release on both HMCs is 6.3.12.0.

Example 6-2 Current IBM Copy Services Manager release

```
dscli> lssoftware -l -type csm -hmc all
Date/Time: September 29, 2025 at 1:49:09 PM EDT IBM DSCLI Version: 7.10.10.348 DS:
IBM.2107-78NKG41
Type Version                  Status  HMC
=======================================
 CSM V6.3.12.0-a20240620-1649 Running 1
 CSM V6.3.12.0-a20240620-1649 Running 2
dscli>
```

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 to be installed.

1. On the IBM Fix Central page, select **IBM Copy Services Manager** as the product, **6.3.15.0** as the installed version, and **Linux PPC** as the platform. Figure 6-5 on page 175 shows a summary of selected options.



Figure 6-5 Selected IBM Copy Services Manager Version for HMC

Note: The HMC OS is Linux PPC. Ensure that the correct platform is selected.

2. Be sure to download the correct Linux-ppcle release. Figure 6-6 shows the correct package type selected. Check the Release Notes, and if there is a newer fix pack file, you can use it instead.

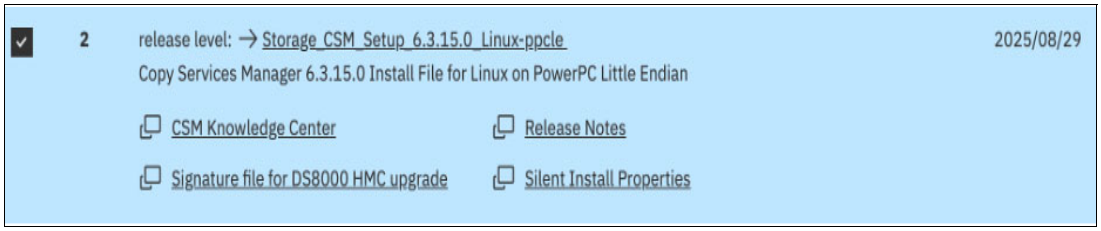


Figure 6-6 IBM Copy Services Manager Linux-ppcle package

When the download process is complete, note the folder path where the files were stored. In the example, the files are stored in the folder `C:\Downloads\CSM_Linux`, as shown in Figure 6-7.

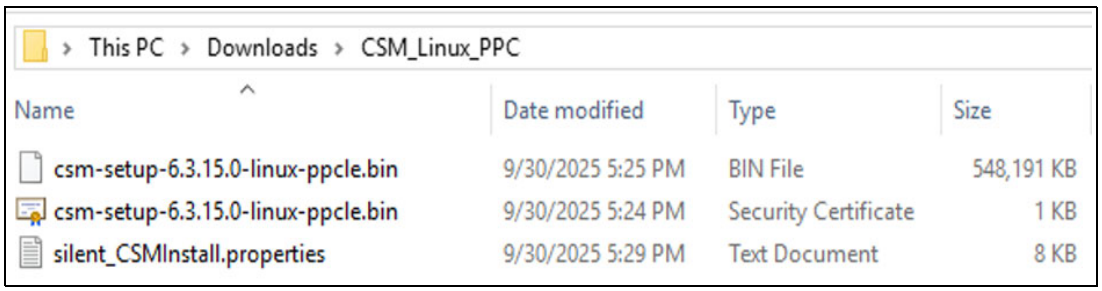


Figure 6-7 Downloaded IBM Copy Services Manager files

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, to the same code level. Failure to do so results in the inability to reconnect Active and Standby CSM servers.

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 6-3 describes the parameters that are necessary for the **`installsoftware`** command.

Table 6-3 DS CLI installsoftware parameters

Parameter	Explanation
`-type csm`	The software type of the installation package.
`-loc software_package`	The full path of the installation package to be installed.
`-certloc certificate_location`	The full path of the certificate file location.
`-hmc 1 | 2 | all`	Specifies the primary or secondary HMC where the software is to be installed. The default is `all`.

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-6.3.15.0-linux-ppcle.bin -certloc
C:\Downloads\CSM_Linux\csm-setup-6.3.15.0-linux-ppcle.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 V6.3.15.0-a20250818-1146 Running
 CSM V6.3.12.0-a20240620-1649 Running
```

The next step is to update IBM Copy Services Manager on HMC2, as shown in Example 6-4 on page 177.

Example 6-4 IBM Copy Services Manager update on HMC2

```
dscli> installsoftware -type csm -loc
C:\Downloads\CSM_Linux\csm-setup-6.3.15.0-linux-ppcle.bin -certloc
C:\Downloads\CSM_Linux\csm-setup-6.3.15.0-linux-ppcle.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 V6.3.15.0-a20250818-1146 Running
 CSM V6.3.15.0-a20250818-1146 Running
```

The IBM Copy Services Manager upgrade is now complete.

6.3 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 DS8A00. It provides remote access to system utilities.

6.3.1 Logging in to the HMC WUI

To log in to the WUI, complete the following steps:

1. Start the Storage Management GUI, as shown in Figure 6-8 on page 178. Click the **Service** icon (wrench) to access the Service Management Console (SMC).
2. Click **Log on and launch the Hardware Management Console web application** to open the login window, as shown in Figure 6-9 on page 178, and log in. The default user ID is `customer` and the default password is `cust0mer`.

> **Important:** Change the default password. The user credentials for accessing the SMC 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 6.6.3, “Service Management Console User Management” on page 187.

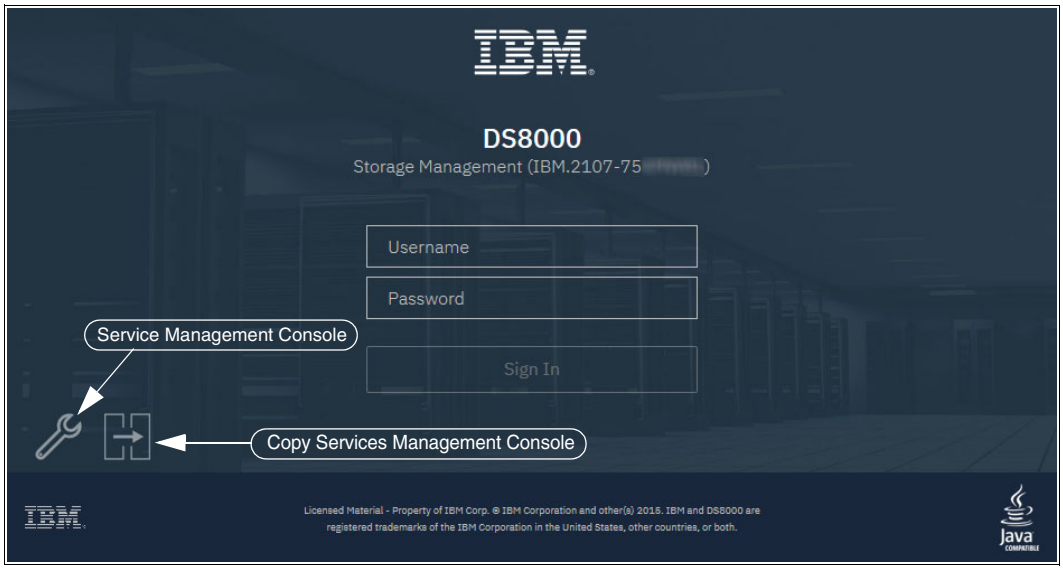


Figure 6-8 DS Storage Management GUI Logon window

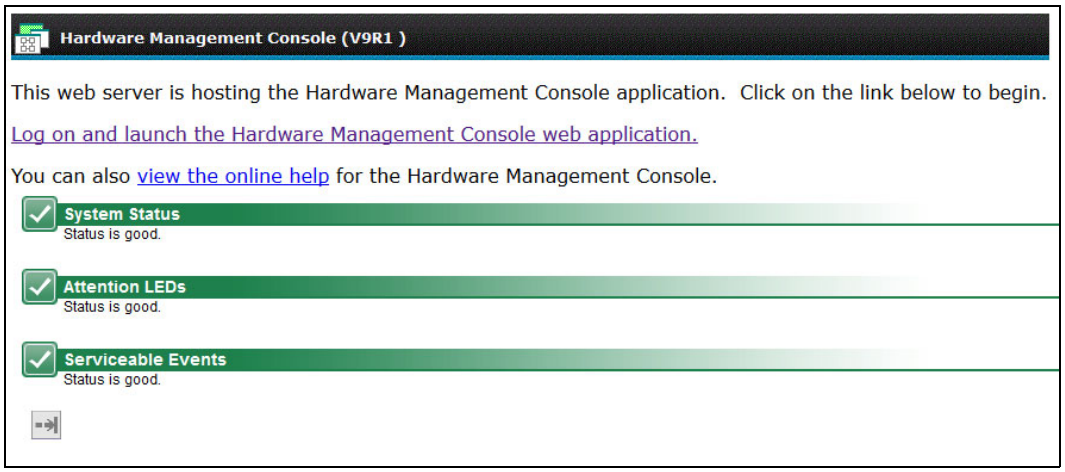


Figure 6-9 Service Management Console application

3. If you are successfully logged in, you see the HMC window, in which you can select **Status Overview** to see the status of the DS8A00. Other areas of interest are shown in Figure 6-10 on page 179.

Because the HMC web UI is mainly a services interface used by IBM service engineers, it is not covered here. For more information, see the **Help** menu.

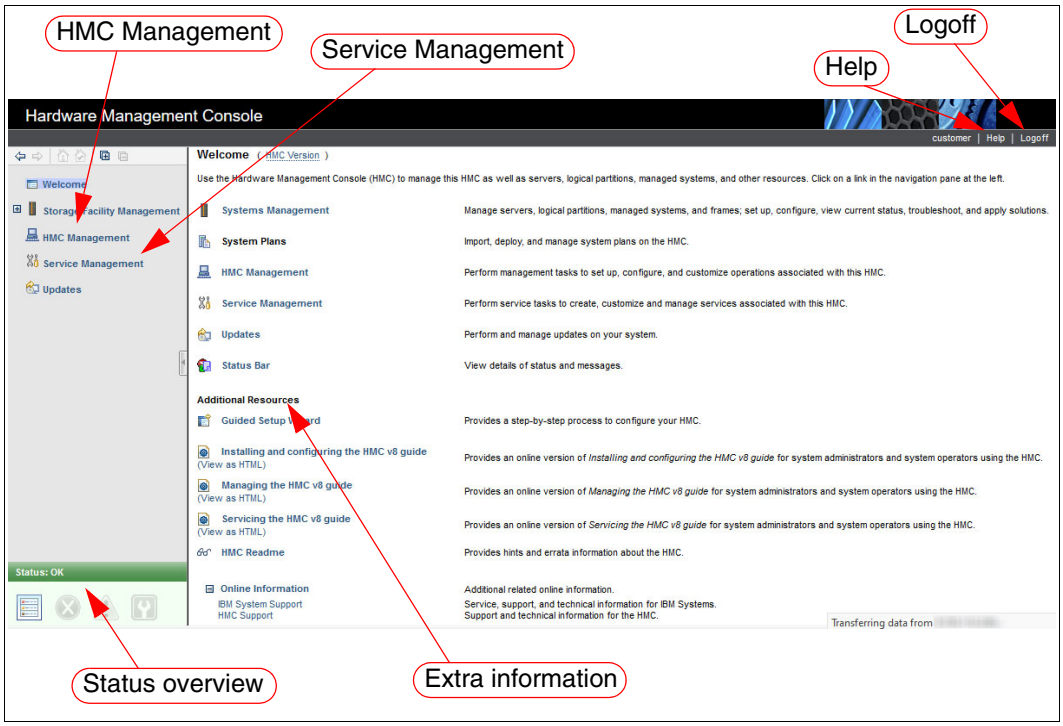


Figure 6-10 Web UI main window

6.3.2 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 DS8A00. It is also referred to as the DS Network Interface (DSNI).

6.4 Management Console activities

This section covers planning and maintenance tasks for the DS8A00 HMC. For more information about overall planning, see Chapter 5, “IBM DS8000 physical planning and installation” on page 143.

6.4.1 Management Console planning tasks

To plan the installation or configuration of the HMC, you must do the following tasks:

- A connection to the client network is needed at the base rack for both HMCs. The connections must be standard CAT5/6 Ethernet cabling with RJ45 connectors.
- IP addresses for the primary and secondary HMCs are needed. The DS8A00 can work with IPv4 and IPv6 networks. From DS8000 release 10 each HMC is capable of supporting two ethernet connections for the customer network. For more information about procedures to configure the DS8A00 HMC for IPv6, see 6.5, “Management Console network settings” on page 182.
- 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 Storage DS8000 (10th Generation) 10.1.0 Introduction and Planning Guide, G10-I53-01.

- The IP addresses of Simple Network Management Protocol (SNMP) recipients must be identified if the client wants the DS8A00 HMC to send traps to a monitoring station.
- Email accounts must be identified if the client wants the DS8A00 HMC to send email messages for problem conditions.
- The IP addresses of Network Time Protocol (NTP) servers must be identified if the client wants the DS8A00 HMC to use NTP for time synchronization.
- When a DS8A00 is ordered, the license and certain optional features must be activated as part of the customization of the DS8A00. For more information, see Chapter 7, “IBM DS8000 G10 features and licensed functions” on page 195.
- The installation tasks for the optional external HMC 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.

6.4.2 Planning for Licensed Internal Code upgrades

The following tasks must be considered regarding the LIC upgrades on the DS8A00:

- LIC changes

  IBM periodically releases changes to the DS8A00 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:

  – Onsite 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.

  DS8A00 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 DS8A00, and covers all versions of OS, multipathing software, and firmware. This check is necessary to ensure that the DS8A00 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 representative.

Never proceed with a LIC update without adhering to all prerequisites.

- Maintenance windows

  The LIC update of the DS8A00 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.

6.4.3 Time synchronization

With the DS8A00, the HMC can use the NTP service. Clients can specify NTP servers on their internal or external network to provide the time to the HMC. It is the client's responsibility to ensure that the NTP servers are working, stable, and accurate. An IBM SSR enables the HMC to use NTP servers (ideally at the time of the initial DS8A00 installation). Changes can be made by the client by using the **Change Date and Time** action under **HMC Management** on the HMC WUI.

Important: For correct error analysis, the date and time information must be synchronized on all components in the DS8A00 environment. These components include the DS8A00 HMC, the attached hosts, IBM Spectrum Control, and DS CLI workstations.

6.4.4 Monitoring DS8A00 with the Management Console

A client can receive notifications from the HMC 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 DS8A00 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 DS8A00 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 DS8A00. For more information about the DS8A00 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 HMC 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 DS8A00 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 HMC, but from the DS8A00 through the FICON connections that run between the server and the DS8A00.

6.4.5 Event notification through syslog

To meet ever increasing security requirements, the DS8A00 supports security and logging events that are forwarded to a syslog server. 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.

6.4.6 Call Home and remote support

The HMC uses outbound (Call Home) and inbound (remote service) support.

Call Home is the capability of the HMC to contact the IBM Support Center to report a serviceable event. Remote support is the capability of the IBM SSR to connect to the HMC to perform service tasks remotely. If the IBM Support Center can connect to the HMC to perform service tasks remotely based on the setup of the client's environment, an IBM SSR can connect to the HMC 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.

6.5 Management Console network settings

The DS8A00 HMC 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 6-11 on page 183.

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.

IMPORTANT: DNS resolution of esupport.ibm.com, either by the HMC or via a proxy server is required for call home. Direct use of IP addresses is no longer supported.

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 6-13 on page 184.

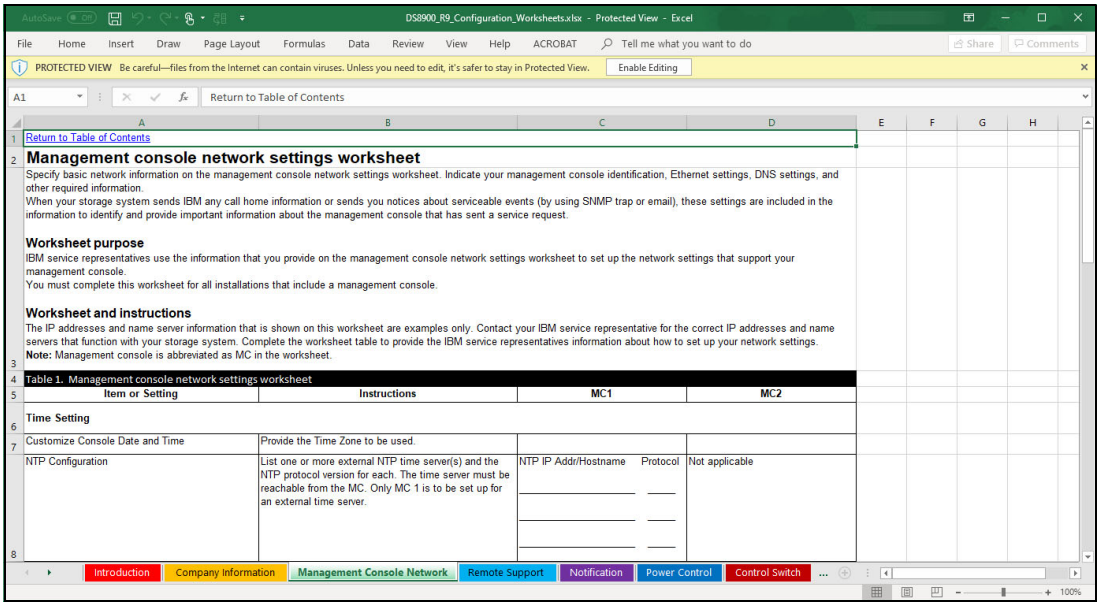


Figure 6-11 DS8000 Configuration Worksheets

6.5.1 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 6-12.

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.

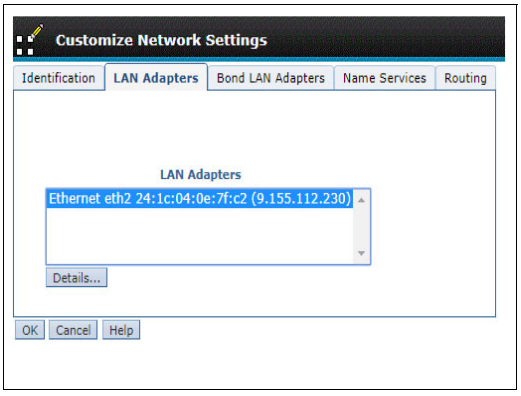


Figure 6-12 Management Ethernet settings

6.5.2 Private networks

The internal private networks (black and gray) are using the IP ranges 172.*x.x.x* by default. 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 HMC 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 6-4 shows the possible options that can be chosen during installation.

Table 6-4 Private networks

Setting	Black network (HMC eht0)	Gray network (HMC 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

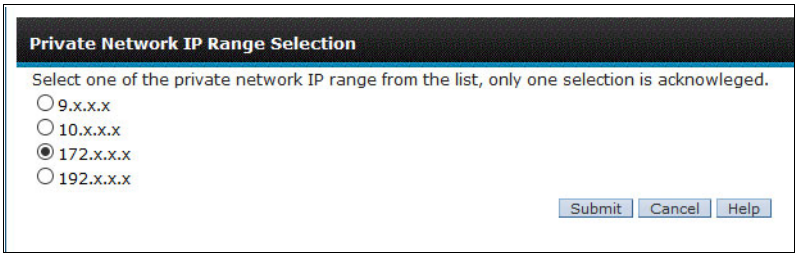


Figure 6-13 Setting the private network range

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.

6.6 User management

The Service Management WUI 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 6.6.3, “Service Management Console User Management” on page 187.

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 DS8A00 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.

6.6.1 Password policies

DS8A00 supports different role-based users. For more information about user and role management, see 8.2, "User and role management" on page 222. 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 HMC 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 DS8A00.

- 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 DS8A00.

6.6.2 Remote authentication

You can enable and configure remote authentication to connect to an LDAP repository.

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.

Figure 6-14 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.

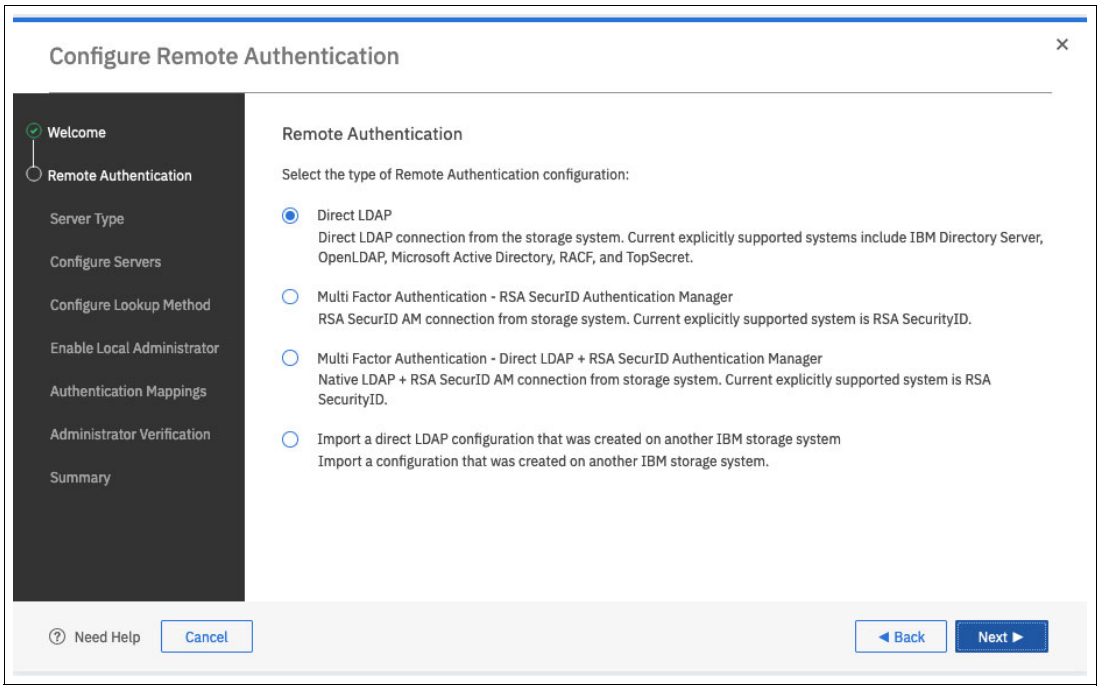


Figure 6-14 Remote Authentication wizard

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).

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.

6.6.3 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 6-15 shows the two options.

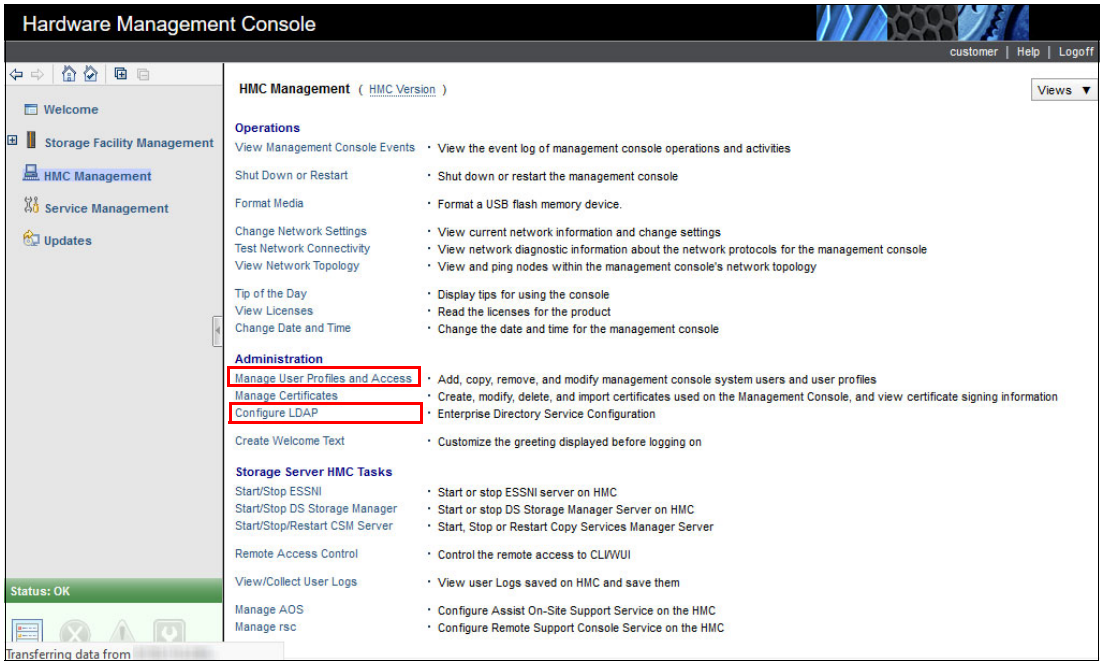


Figure 6-15 Administration options

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 6-6 on page 188.

There are three predefined user roles that are related to the Customer, Service, and Engineering user IDs, as shown in Table 6-5 on page 188.

Table 6-5 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.
esshmcpe	Requires the IBM proprietary challenge/response key for remote access.

The roles, access, and properties for each user ID are described in Table 6-6.

Table 6-6 User roles

Role	esshmccustomer	esshmcserv	esshmcpe
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 6.3, “Web User Interface” on page 177.
2. In the HMC Management window, select **Manage User Profiles and Access**, as shown in Figure 6-16 on page 189.

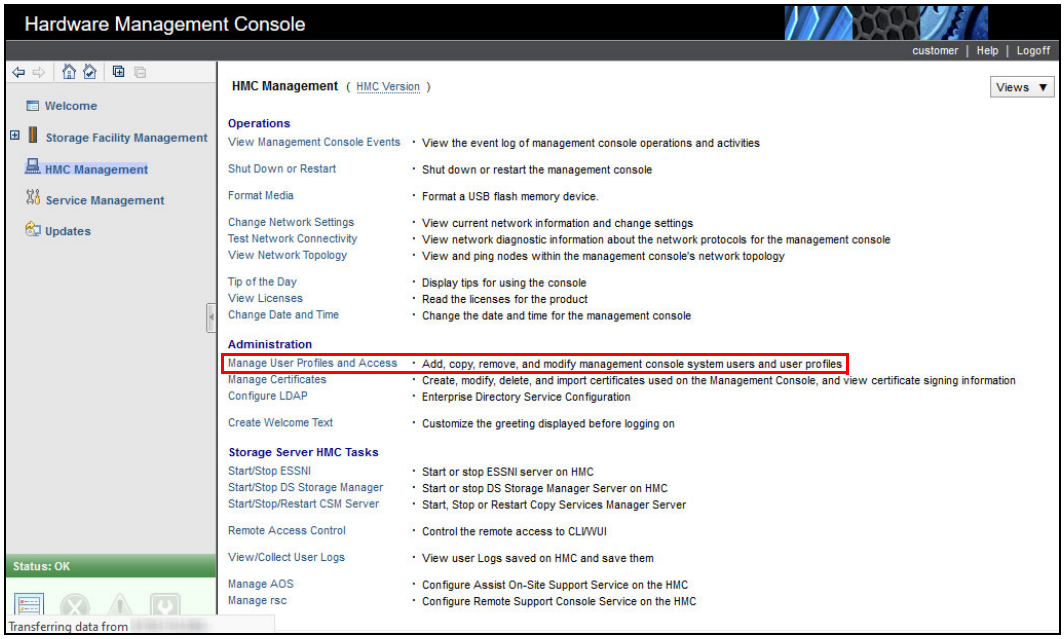


Figure 6-16 Manage User Profiles and Access

A new window opens that lists the user IDs and profiles for the defined console users, as shown in Figure 6-17.



Figure 6-17 Predefined User Profiles

User IDs PE, CE, and customer are specifically for DS8A00 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 `esshmcoperator`.

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 6-18.

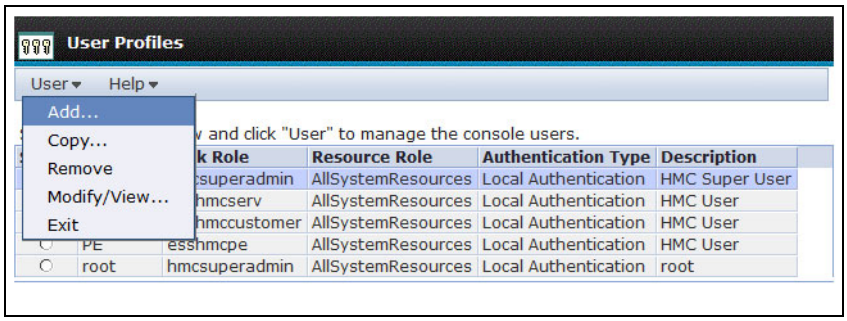


Figure 6-18 User Profiles option menu

2. Click **Add**. The Add User window opens, as shown in Figure 6-19.

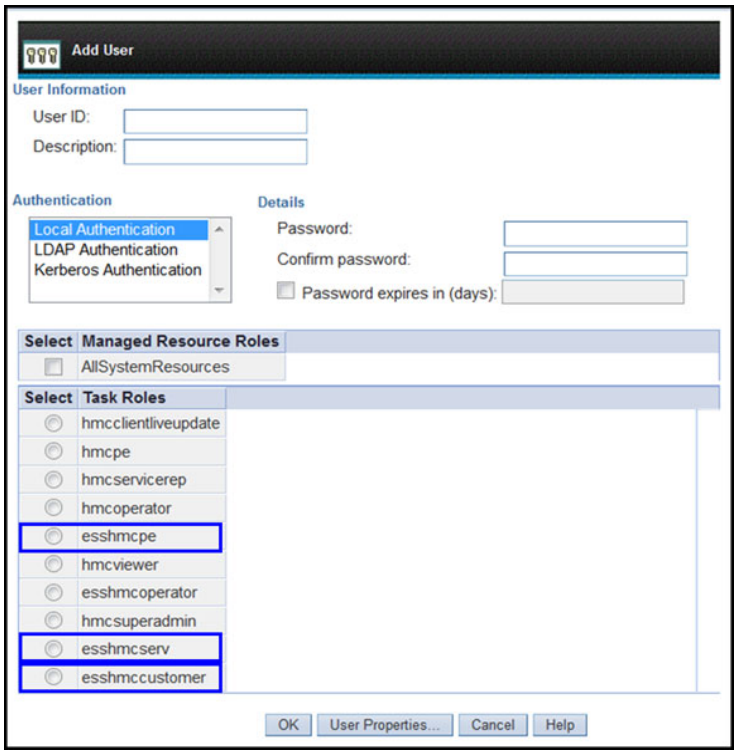


Figure 6-19 Add User window

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 DS8A00 password policies. For more information, see , “After you issue that command, you can run other commands.” on page 185.
   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 6-20 shows these settings.

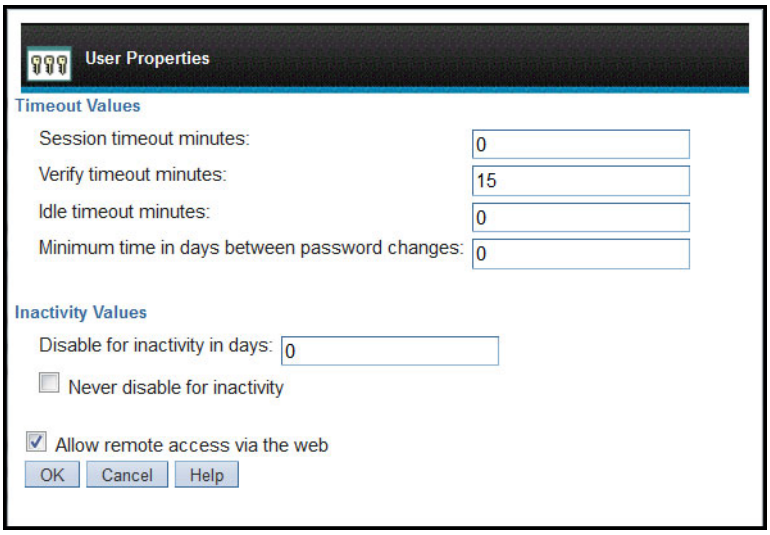


Figure 6-20 User Properties menu

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 6-21 and Figure 6-22 on page 192.

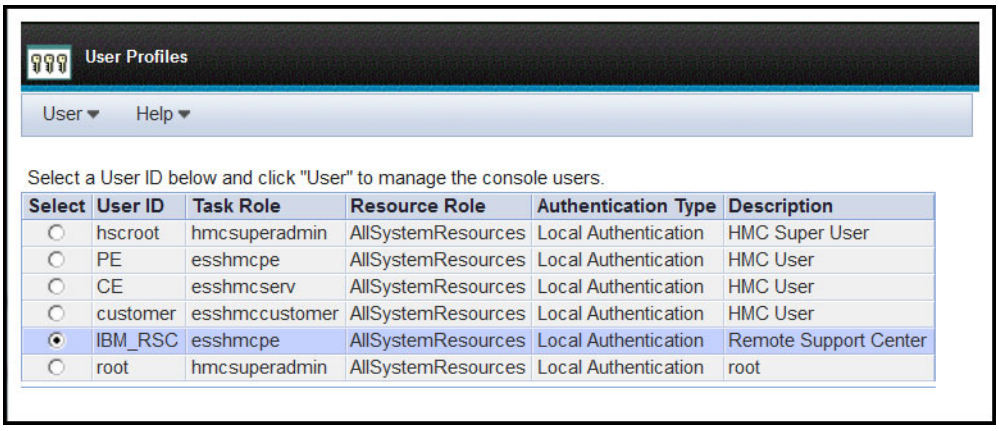


Figure 6-21 IBM_RSC user ID

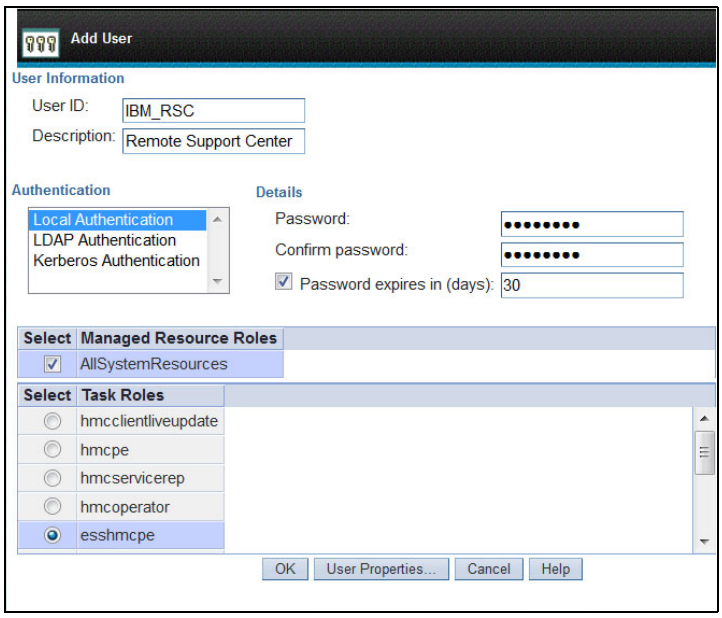


Figure 6-22 IBM_RSC user ID properties

6.6.4 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 6-23.

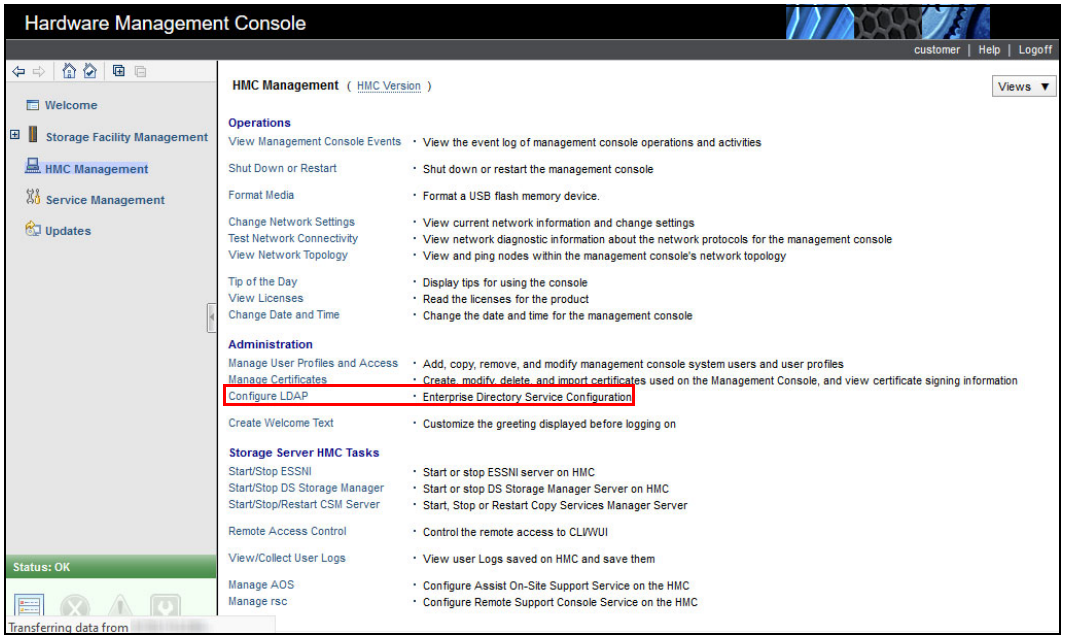


Figure 6-23 Configuring LDAP

2. The window that is shown in Figure 6-24 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.

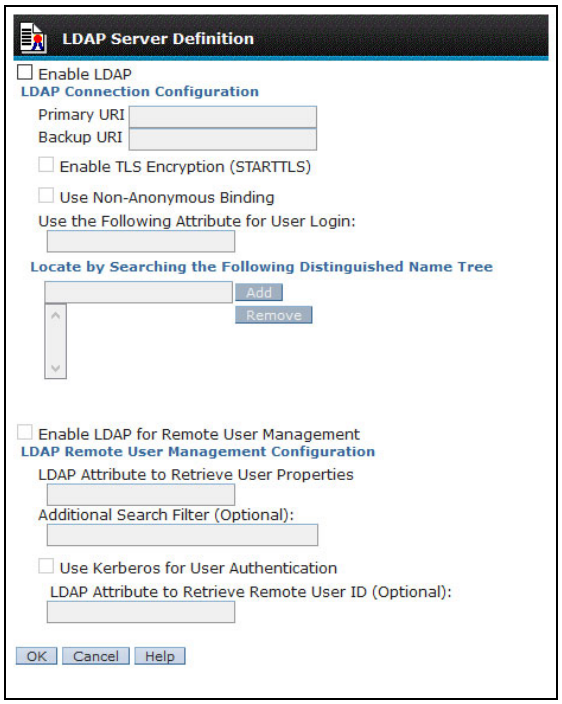


Figure 6-24 LDAP Server Definition window

6.6.5 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.

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).

The RACF PassTicket is a one-time password which can be used within an LDAP environment. Communication between NIClient-ESSNI, and ESSNI-RACF remains untouched.

Administrator must install encryption key using **`managepassticket`** DSCLI Command. Once PassTicket is configured, the user should be able to run inline commands with no pwd parameter specified.

For more information, see *LDAP Authentication for IBM Storage DS8000 Systems: Updated for DS8000 Release 9.3.2*, REDP-5460.

For more information, see IBM DS8000 (10th Generation) Series 10.1.0 Command-Line Interface User's Guide.

6.7 Secondary Management Console

The secondary HMC is used for redundancy and it is part of the DS8A00. The primary HMC is referred to as *HMC1*, and the secondary HMC is referred to as *HMC2*. The two HMCs run in a dual-active configuration, so either HMC can be used at any time. Each HMC is assigned a role of either primary (normally HMC1) or secondary (normally HMC2). Certain service functions can be performed only on the primary HMC.

Important: The primary and secondary HMCs are not available to be used as general-purpose computing resources.

6.7.1 Management Console redundancy benefits

HMC redundancy provides the following advantages:

- Enhanced maintenance capability

  Because the HMC is the only interface that is available for service personnel, an alternative HMC provides maintenance operational capabilities if the internal HMC fails.

- Greater availability for power management

  Using the HMC is the only way to safely power on or power off the DS8A00. The secondary HMC is necessary to shut down the DS8A00 if the primary HMC fails.

- Greater availability of encryption deadlock recovery

  If the DS8A00 is configured for Full Disk Encryption (FDE) and an encryption deadlock situation occurs, the use of the HMC is the only way to input a recovery key to allow the DS8A00 to become operational.

- Greater availability for Advanced CS

  Because all CS functions are driven by the HMC, any environment that uses Advanced CS must include dual HMCs for operational continuity.

- Greater availability for configuration operations

  All configuration commands must go through the HMC. This requirement is true whether access is through IBM Spectrum Control, DS CLI, or DS GUI. The secondary HMC allows these operations to continue if the primary HMC fails.

When a configuration or CS command is run, the DS CLI or DS GUI sends the command to the first HMC. If the first HMC is unavailable, it automatically sends the command to the second HMC instead. Typically, you do not need to reissue the command.

Any changes that are made by using one HMC are instantly reflected in the other HMC. No host data is cached within the HMC, so no cache coherency issues occur.

7

IBM DS8000 G10 features and licensed functions

This chapter describes licensed functions and their activation for the IBM DS8000 G10.

This chapter covers the following topics:

- “DS8000 licensed functions” on page 196
- “Activating licensed functions” on page 203
- “Licensed scope considerations” on page 217
- “Expert Care and ServicePacs” on page 218

© Copyright IBM Corp. 2025. All rights reserved.

7.1 DS8000 licensed functions

The DS8000 G10 family with its DS8A00 models introduces the following Licensed Functions:

- Base Functions License (feature numbers BAA*x*)
- Copy Services (CS) License (feature numbers BAB*x*)
- z-synergy Services (zsS) License (feature numbers BAC*x*)
- Transparent Cloud Tiering (TCT) Function License (feature numbers BAD*x*)
- Security Function Group (feature numbers BAE*x*)

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-FFA indicator Feature Code numbers
Base Function (BF)	ALL, or Fixed-Block (FB)	BAA0 - BAA6
Copy Services (CS)	ALL, Fixed-Block (FB), or Count Key Data (CKD)	BAB0 - BAB7
Z Synergy Services (zsS)	CKD	BAC0 - BAC8
TCT	CKD	BAD0 - BAD8
Security Function Group	ALL	BAE0 - BAEZ (9372-00*x*)

The grouping of licensed functions facilitates ordering. The licenses establish the extent of IBM authorization for the use of the licensed functions on an IBM Storage DS8000.

The five license bundles contain the following functions:

- Base Functions (BF) license:
  - Logical configuration support for Fixed Block (open systems)
  - Operating Environment License (OEL)
  - Thick/Thin-provisioned Volume support
  - Space reclamation (Unmap)
  - Dynamic Volume Expansion
  - FB Compression through FCM
  - Encryption Authorization SSD/FCM
  - Local Key Manager for Data at Rest
  - Easy Tier
  - Secure Boot
  - Multifactor Authorization
  - LDAP
- Copy Services (CS) license:
  - FlashCopy
  - Metro Mirror (MM)

- Global Mirror (GM)
- Global Copy (GC)
- Metro/Global Mirror (MGM)
- Multiple-Target Peer-to-Peer Remote Copy (MT-PPRC)
- Cascaded Incremental Resync
- Safeguarded Copy
- Sub-capacity Licensing
- HyperSwap

IBM Copy Services Manager (CSM) integrated into CS license group
The CS function group includes and enables the CSM on HMC functionality (previously offered under FC 8451) as part of CS, and for all volumes of DS8000 G10 and later generations. This consolidation reduces license complexity and enables all users to use Copy Services Manager for replication management. 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.

- z-synergy Service (zsS) license:
  - Parallel Access Volumes (PAV)
  - HyperPAV, SuperPAV
  - FICON attachment
  - High-Performance FICON for IBM Z (zHPF)
  - zHPF Extended Distance II
  - FICON Dynamic Routing
  - Dynamic Volume Expansion
  - IBM z/OS Distributed Data Backup (zDDB), IBM Sterling® MFT Acceleration with zDDB
  - zHyperWrite
  - zHyperLink
  - z/DS8000 Easy Tier Application
  - GDPS Heat Map Transfer
  - Fibre Channel Endpoint Security (IFCES) – Authorization and Encryption of data in flight (EdiF)
  - CKD Compression through FCM
- Transparent Cloud Tiering (TCT) function license:
  - Multi-Cloud on premises and off premises support
  - Encryption support
  - Compression for TS7700

  With increasing demand and popularity, the Transparent Cloud Tiering function group is created, and customers can benefit from this stand-alone group, which streamlines support on the previously mentioned functions.
- Security Function Group:
  When using IBM FlashCore Modules, built-in ransomware threat detection based on
  - Real-time data pattern analysis

- Machine Learning-based threat detection
- Inline detection

  For systems that started with R10.0 code (prior to September 2025), this function is grandfathered in and can be used after upgrading to R10.1 and higher code levels.

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.

All maintenance and support falls under Expert Care, which defines the support duration (1, 2, 3, 4, or 5 years) and the service level (Advanced or Premium). When purchasing an IBM DS8000 (machine type 5341), the inclusion of Expert Care is mandatory. For more information, see 7.4, “Expert Care and ServicePacs” on page 218.

- All DS8A00 models are sold with a 1-year warranty. This warranty is extended by Expert Care from 2 to 5 years. The machine type 5341 remains in all cases.
- Each license function authorization is associated with a fixed 1-year function authorization 9031-FFA (for BF, CS, zsS, TCT). For the Security Function Group, the license function authorization is associated with a 9372-001/002/003/004/005 authorization model where the last 3 digits indicate the entitled years of service.
- 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 BAC*x* 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`.

Typically, the DS8000 10th-generation licensed functions are enabled, authorized, managed, activated, and enforced based upon the physical capacity contained in a 5341 system. For each group of licensed functions, specific Feature Code numbers indicate the licensed 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 196. 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, like MM, GM or FlashCopy, are possible and often needed.
- 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 all recent DS8000 generations, 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 DS8000 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.
- 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. If this is used on IBM FlashCore Modules, additional ordering of the Security Function Group is recommended.
- As previously noted, IBM Copy Services Manager on HMC offers full IBM Copy Services Manager functions. Only when having the IBM Copy Services Manager manage non-DS8A00 volumes or when for instance running it outside the HMC on an IBM Z host, it would need to be purchased as a separate product and 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 have been 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.
- Transparent Cloud Tiering (TCT), for direct backups by DS8000 to tape, to internal or to a hybrid cloud, is sold and listed as a separate license, just for those IBM Z clients who use this also.
- As a part of the Base Function group, 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 of industry-standard drives. 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*, 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.
- The Drive Encryption Authorization feature provides data encryption by using IBM Full Disk Encryption (FDE) and key managers, such as IBM Security Guardium Key Lifecycle Manager. For external key managers, the key manager must be licensed separately.
- Local Data-at-Rest Encryption is also listed within the DS8000 as a separate subfeature but is free as part of the Base Function group. Please note that for encryption of data-in-flight support or the TCT backup-to-cloud encryption, an external key manager such as GKLM is still required.

Hardware drive compression, as done by the IBM FlashCore Modules (FCM), is not listed as any sublicense and comes directly with the respective base function group. This function becomes available as soon as the FCM drive hardware is purchased.

For more information about these features, see the IBM Storage DS8000 10.1 Introduction and Planning Guide, G10-I53-01.

7.1.2 Licensing cost structure

Licensing goes by terabytes (TB, decimal), where each TB equals one trillion bytes. For the five possible license bundles, the following general guidelines apply:

- BF license

  The license must be at least equal to the total raw (physical) 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

  CS goes by provisioned capacity. The license can be equal to, less than, or potentially even more than the total raw capacity of the storage system. You must license the terabytes that you use (provisioned capacity) in copy relationships. However, we limit the maximum CS license necessary to the raw capacity.
- 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 physical capacity of all ranks that will be formatted as CKD ranks.

- TCT license

  This follows the rules as for the zsS license, for all Z clients who additionally use the TCT functions. The raw capacity of all CKD ranks counts in that moment.

- Security Function Group license

  When using these, license the physical raw capacity, as for BF.

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 via IBM Storage Insights and so 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 and a Storage Insights installation will be active to monitor usage.

By using a remote connection for Call Home, and when having Storage Insights, the CS license can be based on the usable capacity of the volumes that will potentially be in CS relationships. For industry-standard flash drives, this amount is often 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 (or compression) is used on the DS8000 with a significant number of CS functions, the CS license needs to be equal only to the total raw physical 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 DS8000 is 100 TB but 200 TB of thin-provisioning volumes are in MM, only a 100 TB of CS license is needed.

When using the IBM FlashCore Modules (FCM), compression is always on. Hence for these, ordering the CS license as per what the nominal drive raw capacity says is usually recommended.

Tip (also for industry-standard drives): If you want to ensure future use of overprovisioning, use the raw capacity of the system for CS to ensure coverage.

For FlashCopy volumes, you must count the source plus target volumes as provisioned capacity. Several examples are shown as follows.

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 use 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 DS8A00 requires a 10 TB CS license, and the target DS8900F 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, an industry-standard drive set of sixteen 3.84 TB drives has a 61.44 TB raw physical capacity. Or an FCM compression tier drive set of sixteen 19.2 TB drives has a 307.2 TB raw physical capacity.

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.

Easy Tier is a license feature that is available at no charge. Therefore, it is configured, by default, with the BF license bundle. Likewise, the Thin Provisioning feature comes with the BF license bundle. And so does the functionality for Database Protection (for open and FB), or for compression (when using FCM).

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 7-2.

Table 7-2 Ordering granularity

Tier	Mini-mum (TB)	Maxi-mum (TB)	Tier TB granula-rity	Feature quantity (max.)	Range delta (TB)	Feature numbers
1	1	100	10	10	100	BAA0/BAB1/BAC2/BAD2/BAET
2	101	250	15	10	150	BAA1/BAB2/BAC3/BAD3/BAEU
3	251	500	25	10	250	BAA2/BAB3/BAC4/BAD4/BAEV
4	501	1250	75	10	750	BAA3/BAB4/BAC5/BAD5/BAEW
5	1251	3000	175	10	1750	BAA4/BAB5/BAC6/BAD6/BAEX
6	3001	6000	300	10	3000	BAA5/BAB6/BAC7/BAD7/BAEY
7	6001	12000	600	12	7200	BAA6/BAB7/BAC8/BAD8/BAEZ

Tip: For more information about the features and considerations when you order DS8000 G10 licensed functions, use this link:

`https://www.ibm.com/docs/en/search/ds8000?type=announcement`

and limit the date range to September 2024 and newer.

IBMers and BPs when looking for more details can go to the IBM Federated Catalog and search by using the keywords *DS8A10*, *DS8A50, DS8A80,* or for the advanced functions *DS8A00* as search terms.

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 IBM 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 IBM Entitled Systems Support (ESS).

Before you connect to the ESS 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 ESS 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 ESS website, you must know the serial number and possibly the 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.
3. 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.

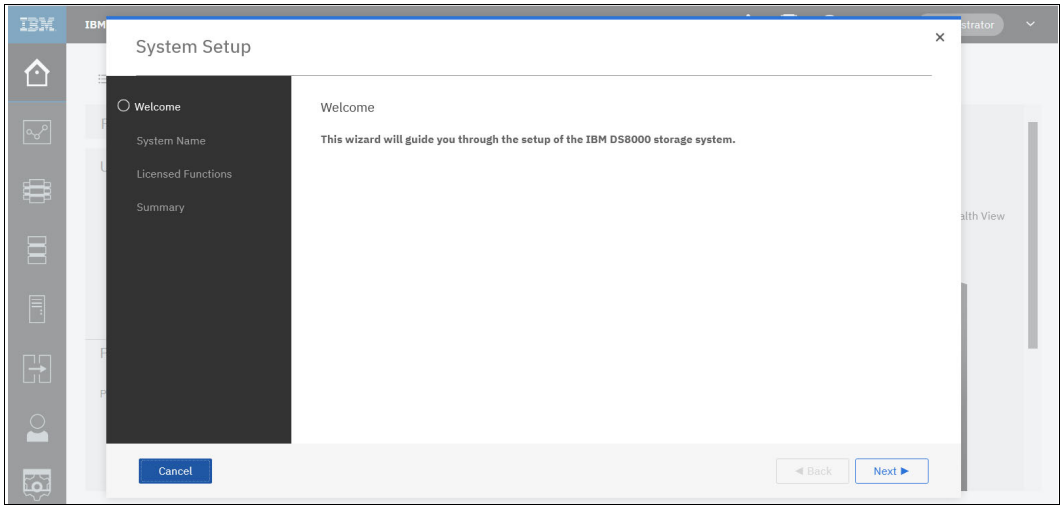


Figure 7-1 System Setup wizard including Licensed Functions activation

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.

4. After an option to give your system another nickname, then to begin the guided procedure to acquire and activate your feature activation keys, select and complete the **Activate Licensed Functions** routine, as shown in Figure 7-2.

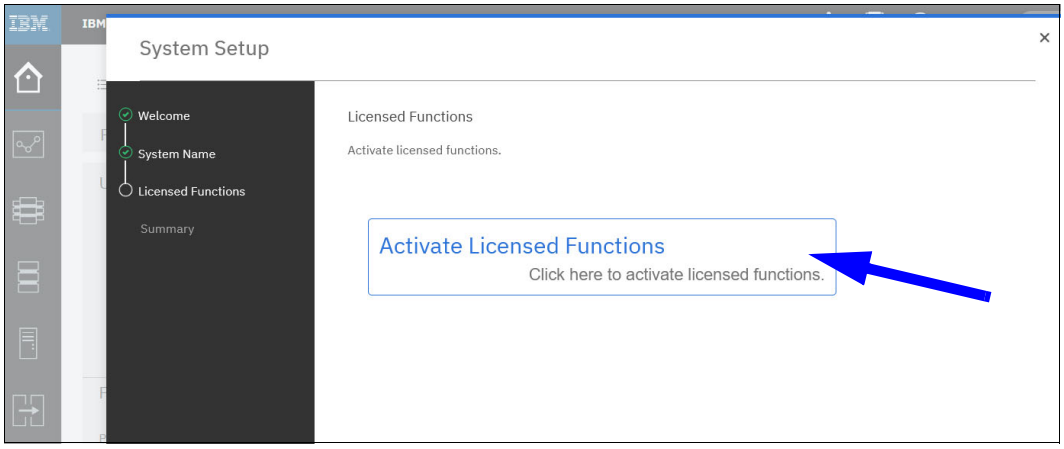


Figure 7-2 Wizard step: Activate Licensed Functions

> **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 ESS website.

5. Use the help function to get further information where and how to obtain the license keys, and how to apply the XML file, as shown in Figure 7-3.

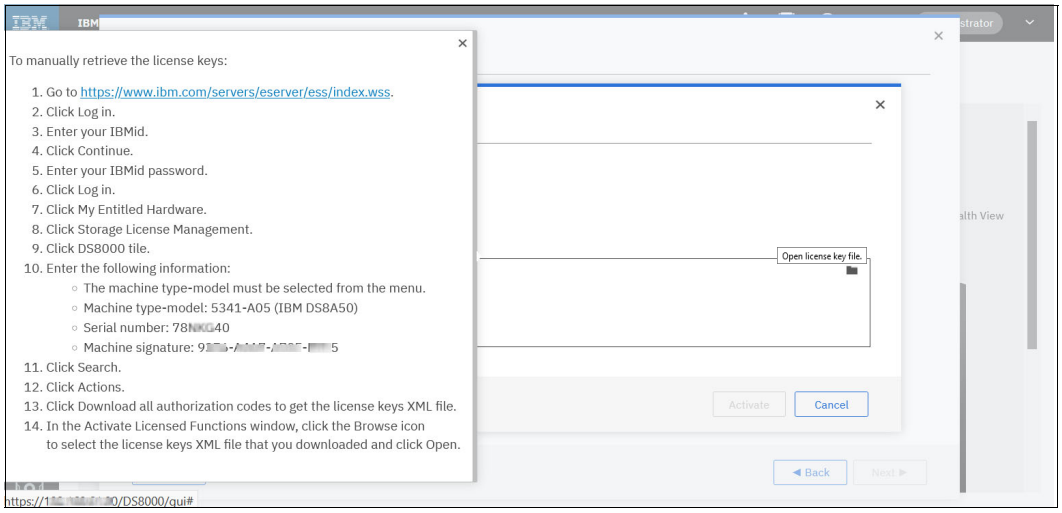


Figure 7-3 Help for adding the license keys

6. All generated license keys become visible after importing, and you click **Activate** to start the activation process, as shown in Figure 7-4 on page 206.

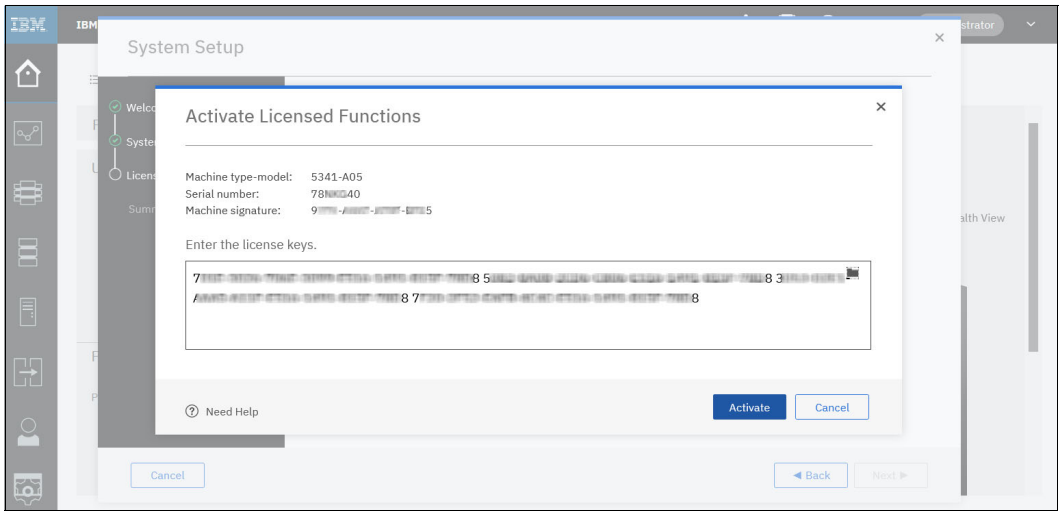


Figure 7-4 Adding the license keys: Activate step

7. The next screens document the activation of the keys, as in Figure 7-5.

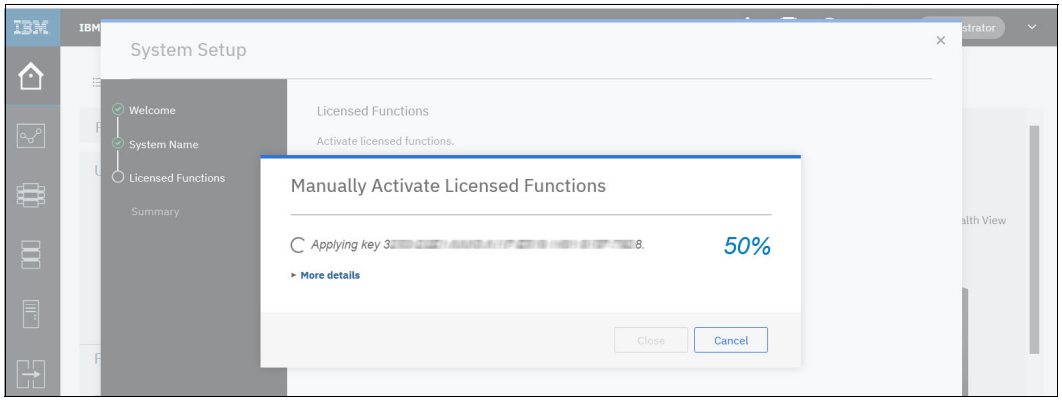


Figure 7-5 License key activation ongoing

8. You can **Activate** further keys manually as needed, or click **Next** to proceed, as in Figure 7-6.

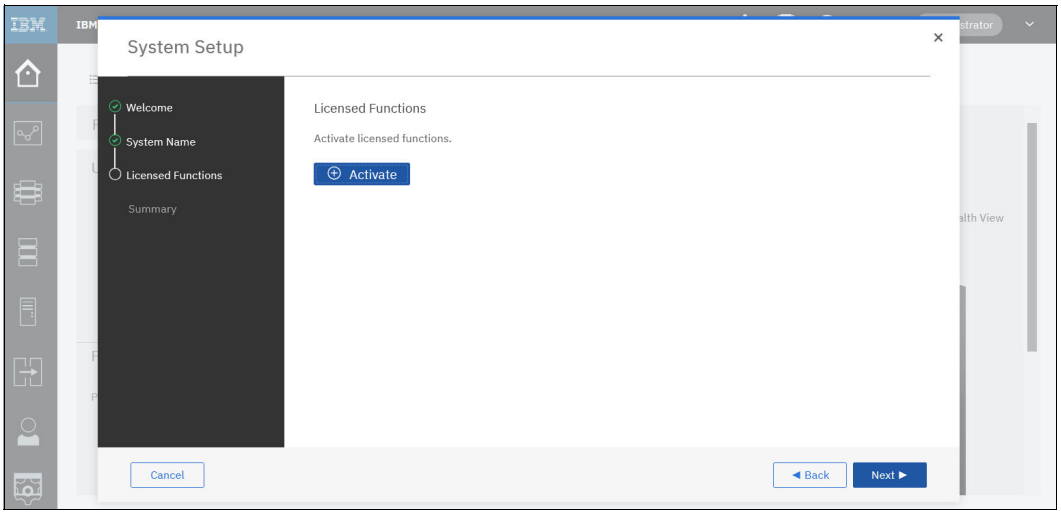


Figure 7-6 Adding the license keys: Further steps

9. 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 7-7. With **Finish** you leave the wizard dialog.

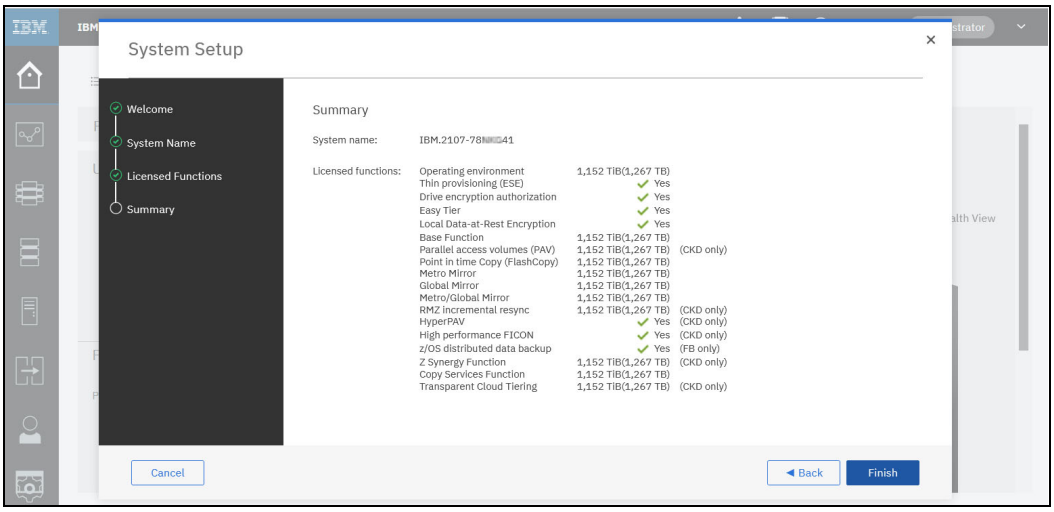


Figure 7-7 Summary of licensed functions

10. 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 7-8.

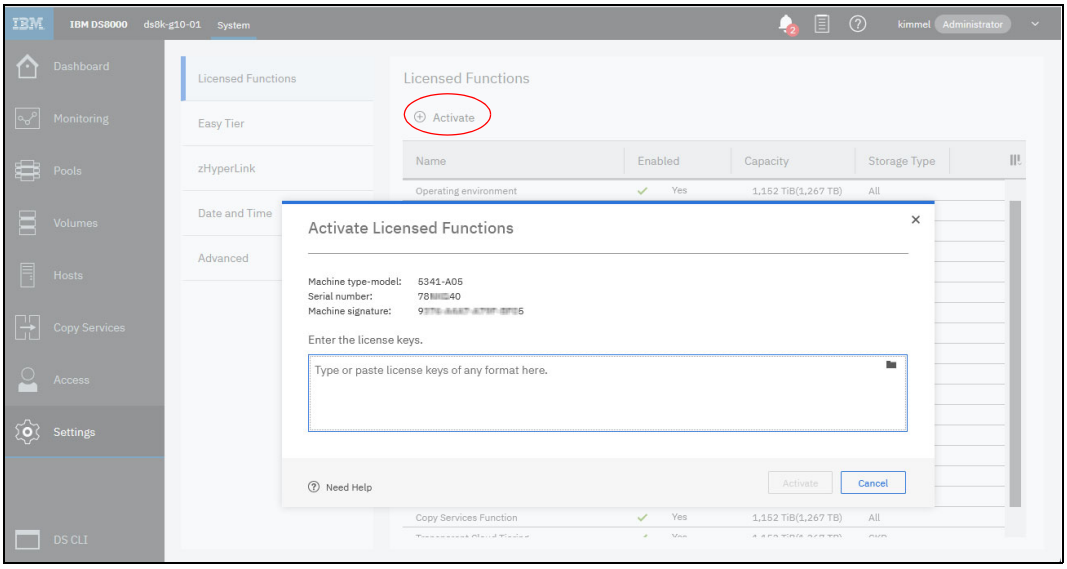


Figure 7-8 Licensed Functions

11. 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 7-9.

Important: The ESS usually expects the Serial Number of the Storage Unit, finishing with a 0, not the one of the Storage Image which finishes with 1.

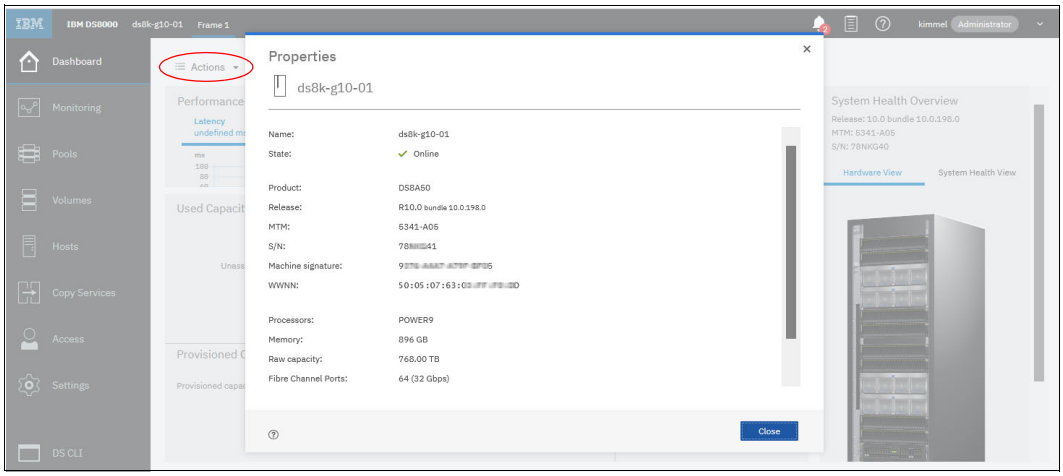


Figure 7-9 Properties window showing the machine signature and MTM

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.

12. Click **Activate** to enter and activate your licensed keys, as shown in Figure 7-4 on page 206.
13. 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 **showsi** commands, as shown in Figure 7-10.

```
dscli> lssi
Name        ID               Storage Unit     Model WWNN             State  ESSNet
==============================================================================
ds8k-g10-01 IBM.2107-78HAL91 IBM.2107-78HAL90 A05   5005076301234567 Online Enabled

dscli> showsi
Date/Time: 27 October 2025 19:56:17 CET IBM DSCLI Version: 7.10.10.348 DS: -
Name             ds8k-g10-01
desc             Saw Shark
ID               IBM.2107-78HAL91
Storage Unit     IBM.2107-78HAL90
Model            A05
WWNN             5005076301234567
Signature        abcd-ef12-3456-7890
State            Online
ESSNet           Enabled
Volume Group     V0
os400Serial      567
NVS Memory       64.0 GB
Cache Memory     832.0 GB
Processor Memory 1015.6 GB
MTS              IBM.5341-78HAL90
numegsupported   16
ETAutoMode       tiered
ETMonitor        automode
IOPMmode         Disabled
ETCCMode         -
ETHMTMode        Enabled
ETSRMode         Enabled
ETTierOrder      High performance
ETAutoModeAccel  Disabled
TDSAutoUpdate    Enabled
TDSSendData      Enabled
DualControl      Disabled
```

Figure 7-10 Obtaining DS8000 information by using the DS CLI

Note: The **showsi** 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 is now mostly `5341`. 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 `A05` for a DS8A50.
- The machine signature, which is found in the Machine signature field and uses the following format:

  `ABCD-EF12-3456-7890`

Table 7-3 documents this information, which is entered at the ESS website to retrieve the activation codes.

Table 7-3 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.

To obtain the activation codes, complete the following steps:

1. As shown in Figure 7-11, connect to the ESS website.

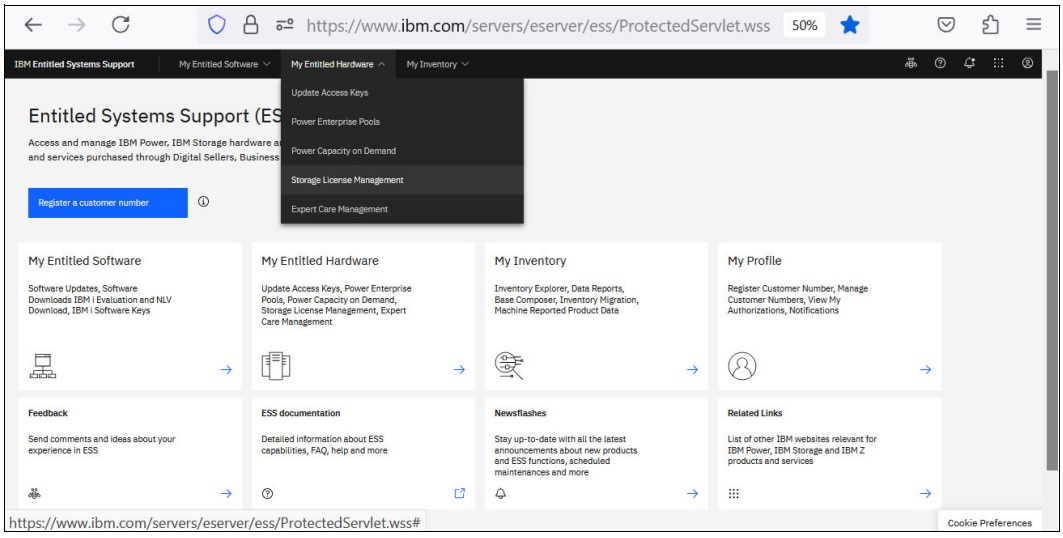


Figure 7-11 ESS website

2. Proceed to **My Entitled Hardware** → **Storage License Management**, as shown in Figure 7-12. Watch out for DS8000.

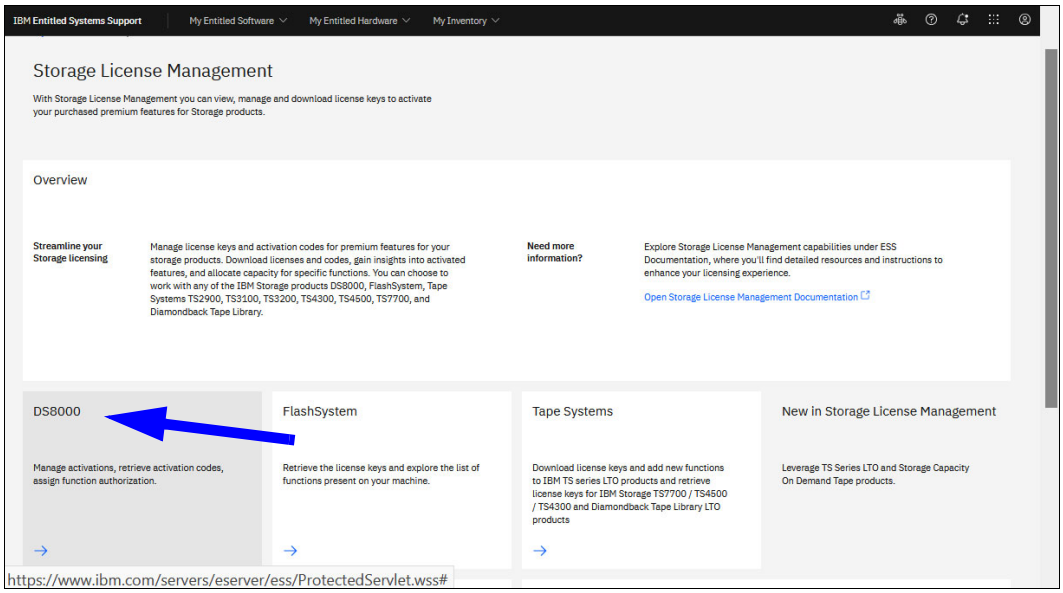


Figure 7-12 ESS Storage License Management

3. Click **DS8000**. The DS8000 machine information window opens, as shown in Figure 7-13.

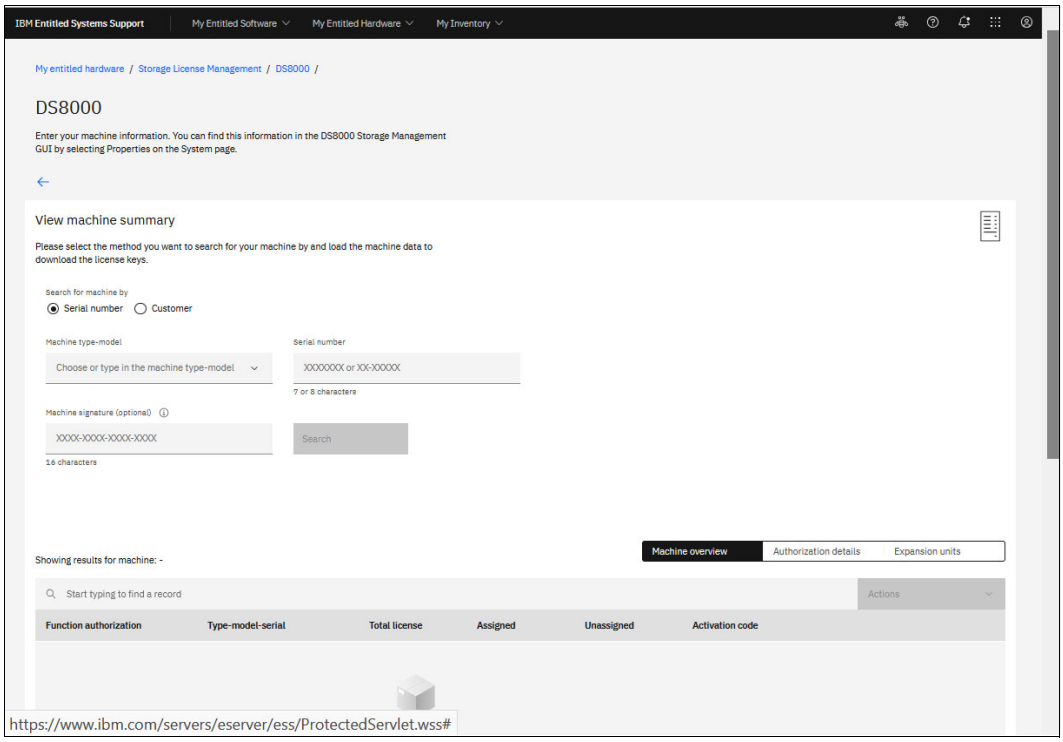


Figure 7-13 ESS DS8000 machine information entry window

4. Enter the machine information that was collected in Table 7-3 on page 211: Look for the model, for instance 5341-A05 (IBM DS8A50), and enter the serial number finishing with 0, as shown in Figure 7-14. Also, the machine signature should be entered when having a brand new machine; only for later viewing it can be optional. Then click **Search**.

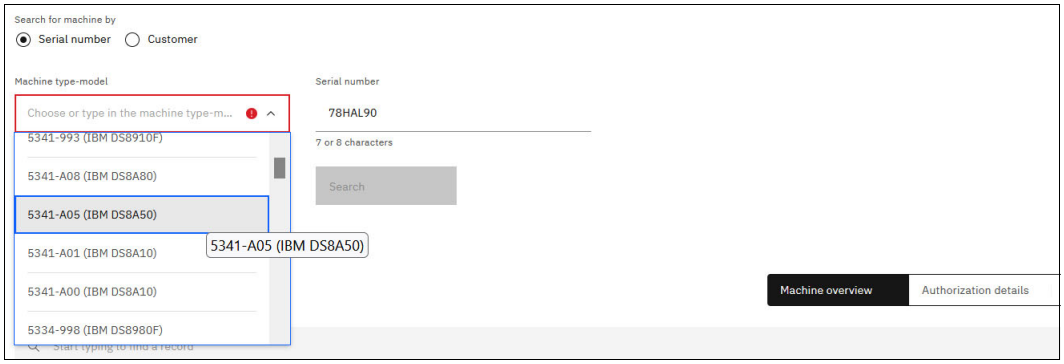


Figure 7-14 DS8000 information: Select model and enter serial number

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`.

5. 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 might show `0.0 TB`, or the results might show a yellow warning message for unassigned license scopes as shown in Figure 7-15.



Figure 7-15 Warning when license scopes are unassigned

6. Figure 7-16 shows a DS8A10 model where licenses are already assigned.

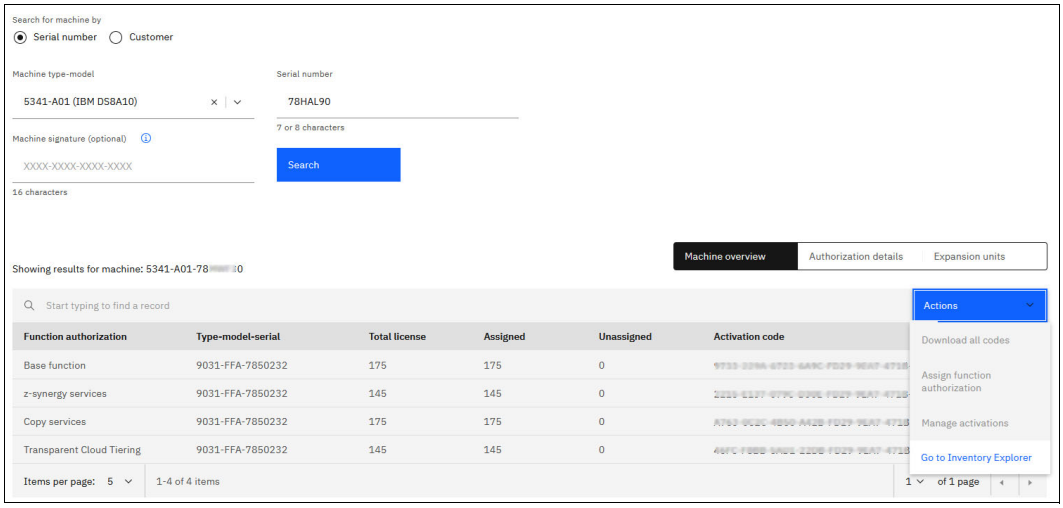


Figure 7-16 DS8000 with assigned licenses

7. Figure 7-17 shows the View machine summary view of a new DS8A10. The yellow warning bar directs you to go to **Actions** → **Manage Activations** to define the correct license scopes.

   This option is only available when also a valid **Machine signature** number has been entered.

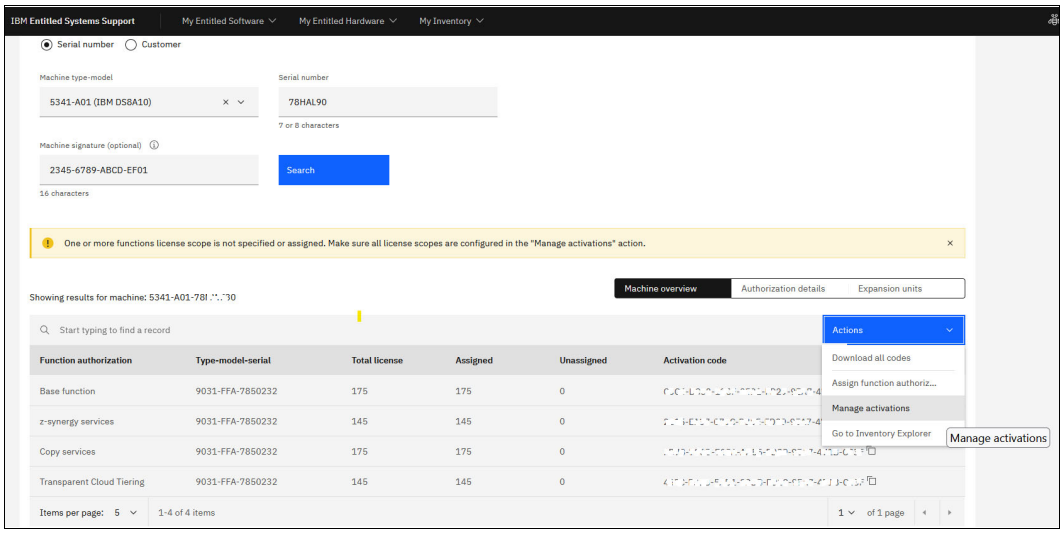


Figure 7-17 ESS View machine summary: New DS8A00 with unassigned scopes

8. If it is not done yet, the `Actions` **Assign function authorization** and **Manage activations** allow you to specify the details for the respective license group. Managing activations is shown in Figure 7-18 on page 215.

   For each license type and storage image, enter the following information:

   – 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.

   – Ensure any **Unassigned** scopes (as shown in Figure 7-18 on page 215) are all assigned. Unassigned scopes can be assigned, for example, to an **All** scope.

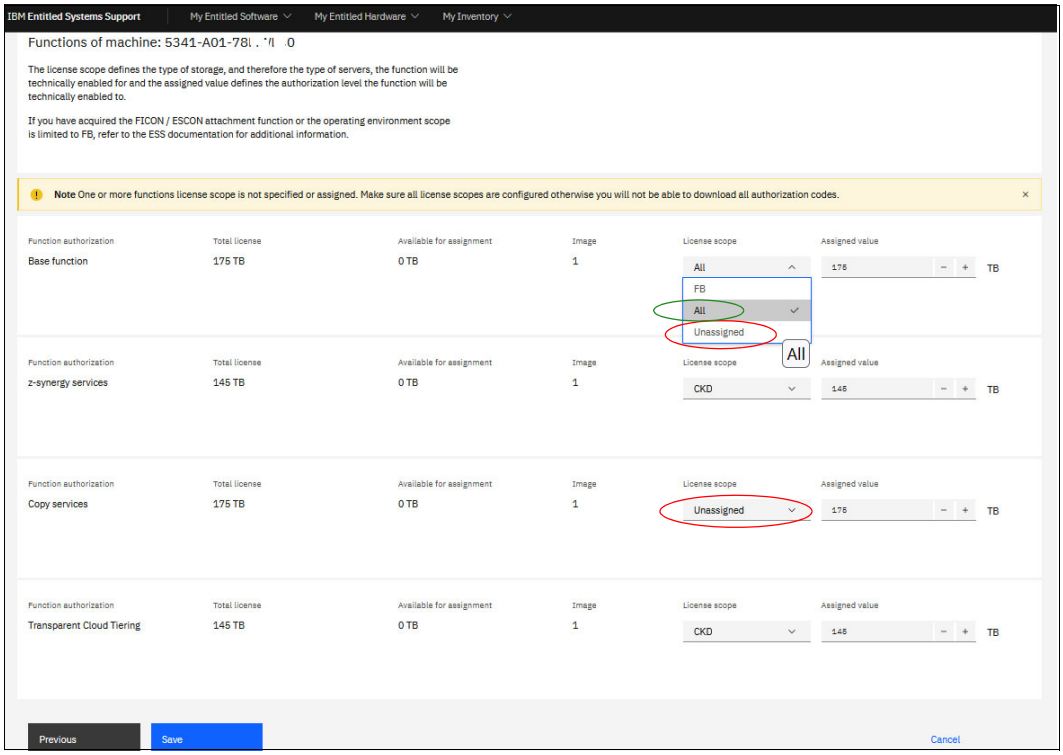


Figure 7-18 ESS: Manage activations

9. Once done, you can either copy each of the generated activation codes, or can proceed to **Download all codes**, what saves the activation codes in an XML file that you can import into the DS8000.
10. This XML file, or the individual keys, can apply in the DS8000 GUI as shown in Figure 7-3 on page 205, and following Figures.

Important: In most situations, the ESS website and application can locate your 9031 licensed function authorization record when you enter the DS8000 5341 serial number, or serial number and signature. However, if the 9031 licensed function authorization record is not attached to the 5341 record, you must assign it to the 5341 record by using the **Assign function authorization** link in the ESS application. In this case, you need the 9031 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 can also 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 **showsi** command to display the DS8000 machine signature, as shown in Figure 7-19.

```
dscli> showsi
Date/Time: 27 October 2025 20:01:45 CET IBM DSCLI Version: 7.10.10.348 DS: -
Name             ds8k-r9-01
desc             Sand Shark
ID               IBM.2107-78HAL91
Storage Unit     IBM.2107-78HAL90
Model            998
WWNN             5005076301234567
Signature        abcd-ef12-3456-7890
State            Online
ESSNet           Enabled
Volume Group     V0
os400Serial      567
NVS Memory       127.5 GB
Cache Memory     4168.4 GB
Processor Memory 4343.4 GB
MTS              IBM.5341-78HAL90
numegsupported   16
ETAutoMode       tiered
ETMonitor        automode
IOPMmode         Disabled
ETCCMode         -
ETHMTMode        Enabled
ETSRMode         Enabled
ETTierOrder      High performance
ETAutoModeAccel  Disabled
TDSAutoUpdate    -
TDSSendData      -
DualControl      -
```

Figure 7-19 DS CLI showsi command

2. Obtain your license activation codes from the IBM ESS website, as described in 7.2.2, “Obtaining the activation codes” on page 211.
3. 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:\5341_78XXXX0.xml IBM.2107-78XXXX1
   ```

   Or you can apply individual keys by running the following command:

   ```
   dscli> applykey -key f190-1234-1234-1234-1234-5678-1234-5678 IBM.2107-78XXXX1
   CMUC00199I applykey: License Machine Code successfully applied to storage image
   IBM.2107-78XXXX1.
   ```

4. Verify that the keys were activated for your storage unit by running the **lskey** command, as shown in Figure 7-20 on page 217.

```
dscli> lskey
Activation Key                                Authorization Level (TB) Scope
==============================================================================
Base function                                 1100                     All
Copy services                                 800                      All
Encryption Authorization                      on                       All
Global mirror (GM)                            800                      All
High Performance FICON for System z (zHPF)    on                       CKD
IBM HyperPAV                                  on                       CKD
IBM System Storage DS8000 Thin Provisioning   on                       All
IBM System Storage Easy Tier                  on                       All
IBM z/OS Distributed Data Backup              on                       FB
Local Data-at-Rest Encryption                 on                       All
Metro mirror (MM)                             800                      All
Metro/Global mirror (MGM)                     800                      All
Operating environment (OEL)                   1100                     All
Parallel access volumes (PAV)                 575                      CKD
Point in time copy (PTC)                      800                      All
RMZ Resync                                    800                      CKD
Transparent Cloud Tiering                     575                      CKD
z-synergy services                            575                      CKD
```

Figure 7-20 Using the lskey command to list the installed licenses

For more information, see CLI commands.

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 and ServicePacs

With Expert Care for IBM DS8000, 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 to 5 years.

Expert Care services

IBM Storage Expert Care includes various services, depending on the selected service level. All services are listed in Table 7-4.

Table 7-4 Expert Care services

Services	Advanced	Premium
Hardware Maintenance (IOR = IBM Onsite 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
Onsite Code Load	Optional	Optional
Media Retention	Optional	Optional

The Technical Account Manager is a 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 12.12, “Using IBM Storage Insights” on page 451.

Expert Care feature codes

The MTMs for Expert Care are 5130-A0*x* for Advanced and 5130-P0*x* for Premium, where *x* indicates the number of years of support. The indicator for the Expert Care Premium software program number is 577*x*-ECP.

Table 7-5 on page 219 shows the Feature Codes for each of the available options.

Table 7-5 Expert Care Feature Codes

Description	Expert Care Advanced 5130 (A01 - A05)	Expert Care Premium 5130 (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
Onsite 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 available for contact and resolution times, including committed fix times of 4, 6, 8, 12, 24, 48, or 72 hours, each associated with a corresponding feature code. For more details, contact your IBM Sales Representative.

Media Retention ServicePac

Optional ServicePac® services for storage can provide retention of defective media, such as hard drives and flash drives, that are being replaced. If the reported problem requires the replacement of media and the client has purchased the Media Retention ServicePac, a replacement will be provided by IBM and the removed defective media can be retained by the client instead of being returned to IBM.

The respective M/T 6664 ServicePac contract involves a 3-year, or 5-year flash drive media retention. You can contact your sales representative for further details.

8

DS8000 G10 configuration flow

This chapter provides a brief overview of the sequence of tasks that are required to configure an IBM DS8000 G10.

This chapter covers the following topics:

- “Configuration worksheets” on page 222
- “User and role management” on page 222
- “Encryption” on page 224
- “Network security” on page 224
- “Configuration flow” on page 225
- “General storage configuration guidelines” on page 226

© Copyright IBM Corp. 2025. All rights reserved.

8.1 Configuration worksheets

Before a new DS8000 G10 system is delivered, it is required to complete the DS8000 G10 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 Storage DS8000 10.1 Introduction and Planning Guide, G10-I53-01. The guide provides detailed information to help you plan a successful installation of the system.

The purpose of the configuration worksheets is to enable a smooth installation of the DS8000 G10 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 DS8000 G10 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 as quickly as possible during a critical event .
- 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, 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 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 service methods, such as the **`list`** and **`show`** commands.
- The IBM Service role (ibm_service) has access to all 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, SG24-8367*.

For more information about the capabilities of certain user roles, see User roles or use the DS GUI Help function.

The DS8000 G10 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. For more information about remote authentication and LDAP for the DS8000 G10, see Using remote authentication.

8.3 Encryption

More planning is required if you intend to activate data-at-rest encryption for the DS8000 G10 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: DS8000 G10 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.

  The DS8000 G10 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 DS8000 G10. 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, DS8000 G10 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 DS8000 G10 system.
- IBM Fibre Channel Endpoint Security: Use this feature to encrypt data as it is transmitted between your IBM Z server and DS8000 G10. This encryption mechanism also uses external key servers. It can be enabled at any time.
- Transparent Cloud Tiering (TCT): The DS8000 G10 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 DS8000 encryption, see 5.3.6, “Key manager servers for encryption” on page 162 and *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security, 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 DS8000 G10 (specifically the HMC) is important depending on the client requirements. The DS8000 G10 supports compliance with the National Institute of Standards and Technology (NIST) SP 800-131a standards.

- DS8A00 now support both Transport Layer Security (TLS) 1.2 and 1.3, 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 DS8000 G10. 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 DS8000 G10 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 DS8000 G10 Storage Management GUI” on page 229.

If you perform the logical configuration by using DS CLI, the following steps provide a high-level overview of the configuration flow. For more detailed information about using and performing the logical configuration with DS CLI, see Chapter 10, “IBM DS8A00 Storage Management Command-line Interface” on page 335.

Here is the general configuration flow:

1. Install license keys: Activate the license keys for the DS8000 G10 storage system. For more information about activating licensed functions, see Chapter 7, “IBM DS8000 G10 features and licensed functions” on page 195.

   Important: If data-at-rest encryption is activated, the encryption configuration must be performed before starting the logical configuration.

2. Create a pool pair: Define the storage type (Fixed Block or Count Key Data). Select the extent unit size. You have the option add arrays to the pool, you can also add storage to the pools in the next step.

   Note: Easy Tier is automatically configured in automatic mode. If this is not desired it can be selected in the custom pool section..

3. Add arrays to the pool: If there are unassigned arrays they can be added to the pool.
4. Consider other controls and monitoring when working with space-efficient volumes. For more information, see *IBM DS8000 Thin Provisioning, REDP-5343*.

   Note: You can modify the Pool capacity (Logical Capacity) and Compression capacity (Physical Capacity) thresholds from their default value of 85%.

5. 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).
6. Create Hosts : Define open systems hosts and their FC host bus adapter (HBA) worldwide port names (WWPNs)
7. Create the open systems volumes: Create striped open systems FB volumes and assign them to one or more hosts.
8. Create the IBM Z logical control units (LCUs): Define their type and other attributes, such as subsystem identifiers (SSIDs).
9. 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 storage 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.
- For best performance use max two poolsper volume type. Best performance is achieved with as many arrays as possible in the pool pair.
- 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 3.4.5, “Logical subsystems” on page 100.
- Volumes of one LCU or LSS can be allocated on multiple storage pools that are assigned to the same internal server.
- The following options are available for FB pools:
  - Create a host for each server.
  - Map the required volumes to the host.
  - If LUN sharing is required, the following items provide two possible use cases (Figure 8-1 on page 226):
    - Application sharing:
      Create one host for each server. Assign the shared volumes to each host.
      The advantage of this option is that you can assign private and shared volumes to each host.
    - Clustering:
      Create one host for each server.
      Create a Cluster and assign both hosts to the cluster.
      Map the volumes to the cluster, the volumes will then be mapped to all hosts in the cluster with the same SCSI id.

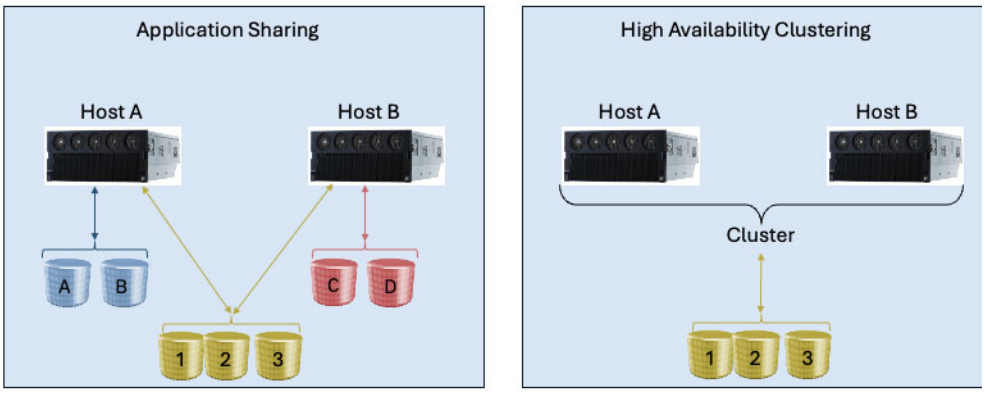


Figure 8-1 Shared volumes

- 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: It is recommended to have dedicated ports for copy services. For more information see DS8000 Host Adapter Configuration Guidelines.

9

IBM DS8000 G10 Storage Management GUI

This chapter describes the DS8000 Storage Management DS GUI that is available with the IBM DS8A00 system. In the DS GUI, you can 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” on page 230
- “DS GUI: Getting started” on page 231
- “Managing and monitoring the storage system” on page 237
- “System configuration overview” on page 245
- “Logical configuration overview” on page 268
- “Logical configuration for open systems volumes” on page 269
- “Logical configuration for Count Key Data volumes” on page 289
- “Expanding volumes” on page 300
- “Deleting a pool” on page 302
- “Deleting volumes” on page 302
- “Reinitializing a thin-provisioned volume” on page 304
- “Easy Tier support” on page 304
- “Monitoring system health” on page 305
- “Performance monitoring” on page 319
- “Fibre Channel error rate statistics” on page 333
- “Providing feedback” on page 334

© Copyright IBM Corp. 2025. All rights reserved.

9.1 Introduction

The DS GUI is packed with usability features that offer an easy-to-use interface for administration of the DS8000.

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. It has been continuously enhanced to deliver a consistent graphical experience, making it easier for users to transition between solutions such as IBM FlashSystem®, IBM Spectrum Virtualize, 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 DS8A00 Storage Management Command-line Interface” on page 335.

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 Management help functions

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 DS8000 Storage Management GUI web browser support and configuration.

On a new storage system, the user must log on as the administrator. The password expires immediately: the user must change the password.

The initial view of the system is shown in Figure 9-1.

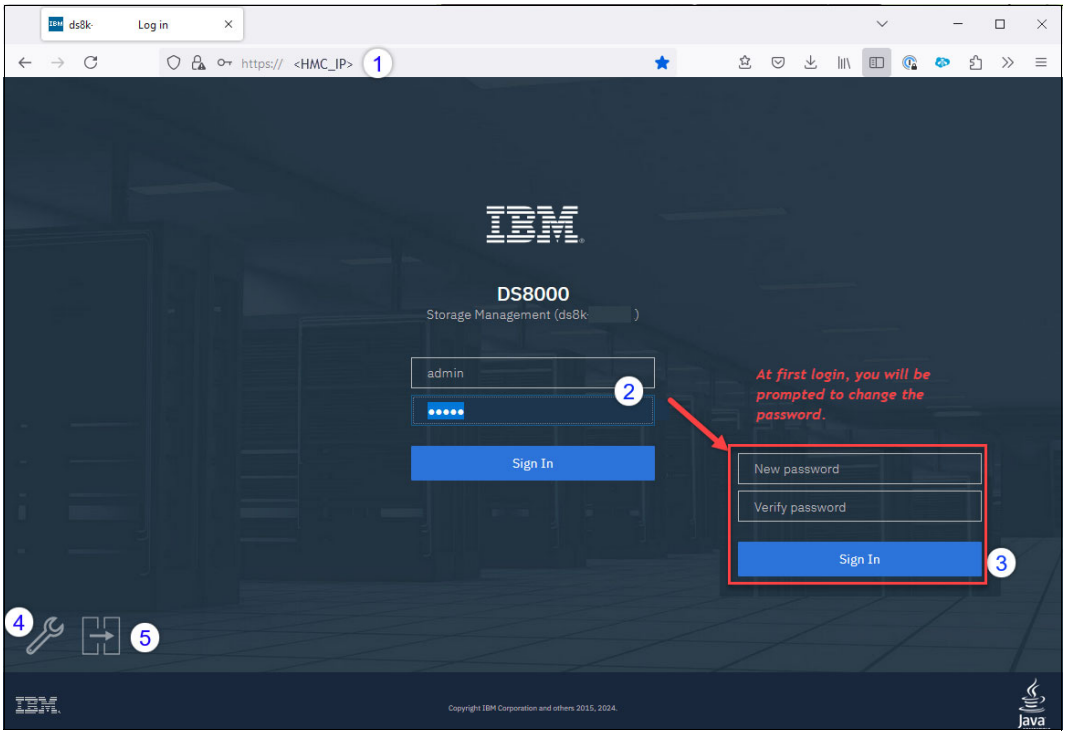


Figure 9-1 Storage Management GUI initial login

Figure 9-1 on page 231 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.

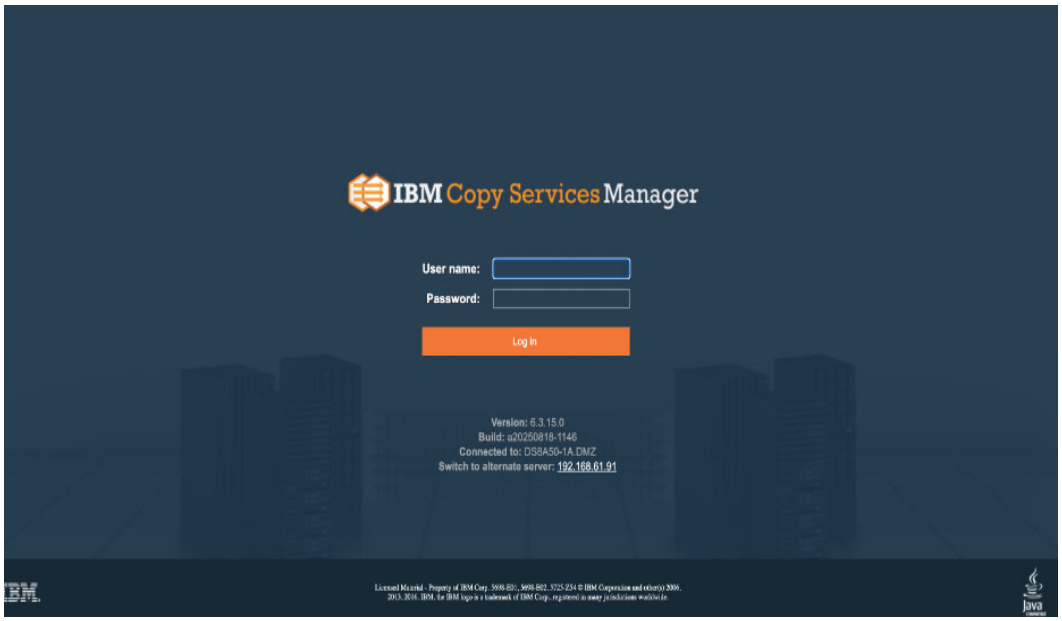


Figure 9-2 IBM Copy Services Manager window started from the DS GUI

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.

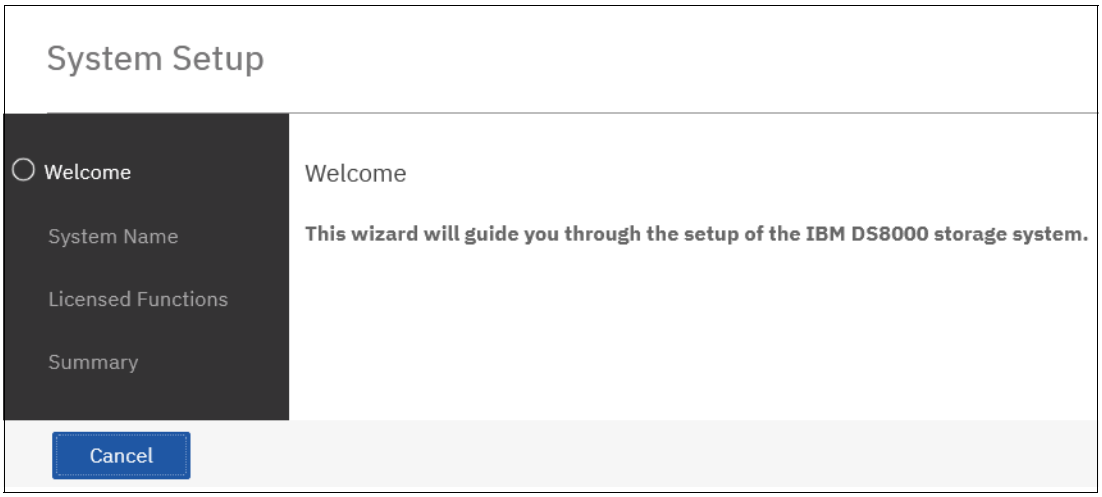


Figure 9-3 System Setup: Welcome pane

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 naming convention that is used in their environment.

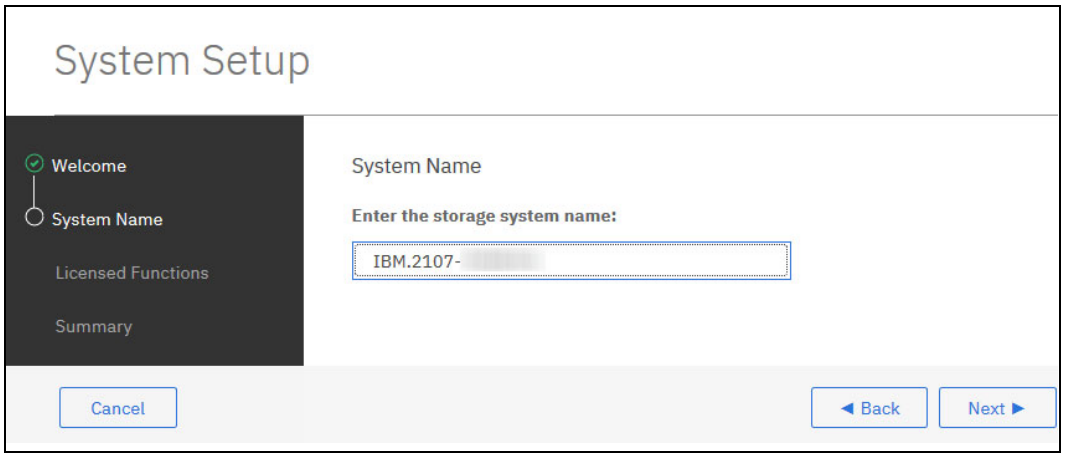


Figure 9-4 System Setup: System Name window

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 IBM Entitled Systems Support (ESS) 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 DS8000 G10 features and licensed functions” on page 195.

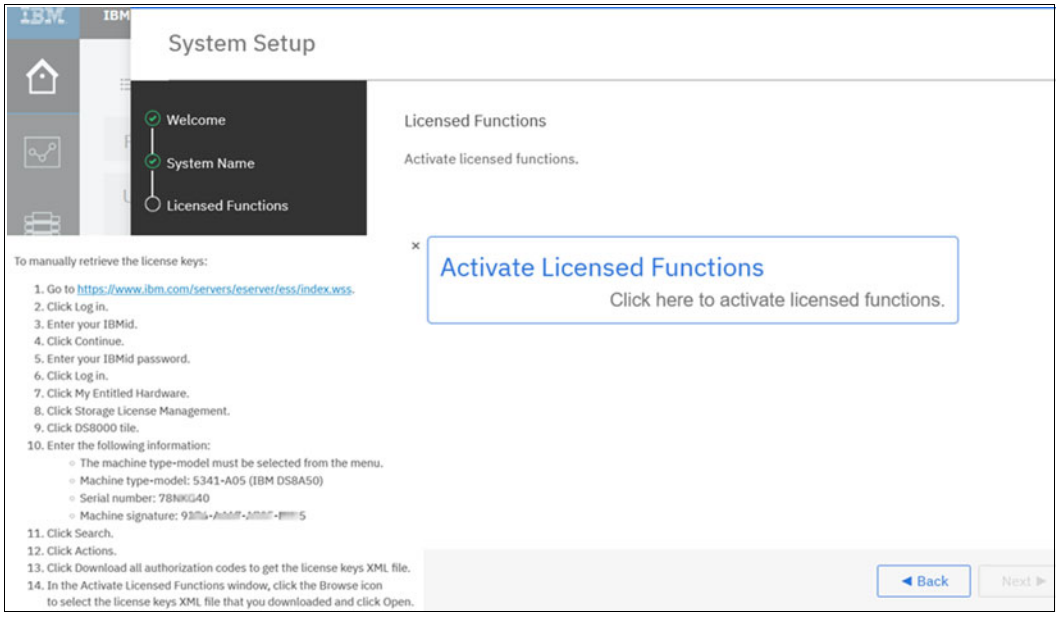


Figure 9-5 System Setup: Activate Licensed Functions

4. When the license keys are entered, click **Activate** to enable the functions, as illustrated in Figure 9-6.

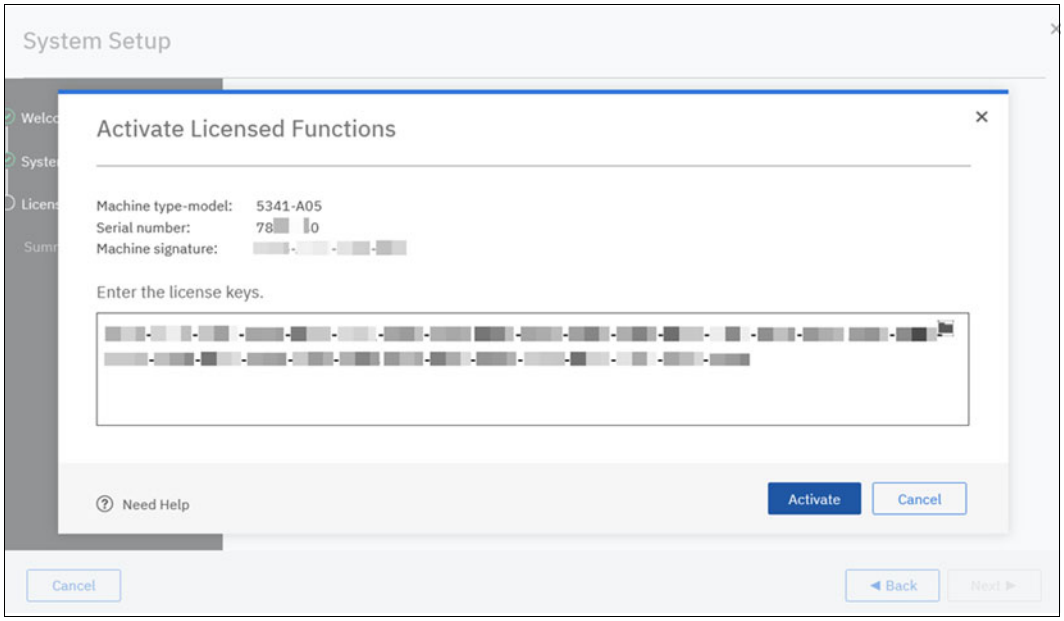


Figure 9-6 System Setup: Enabling Licensed Functions

5. 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 (CSM) on the HMC if the corresponding licenses were activated.

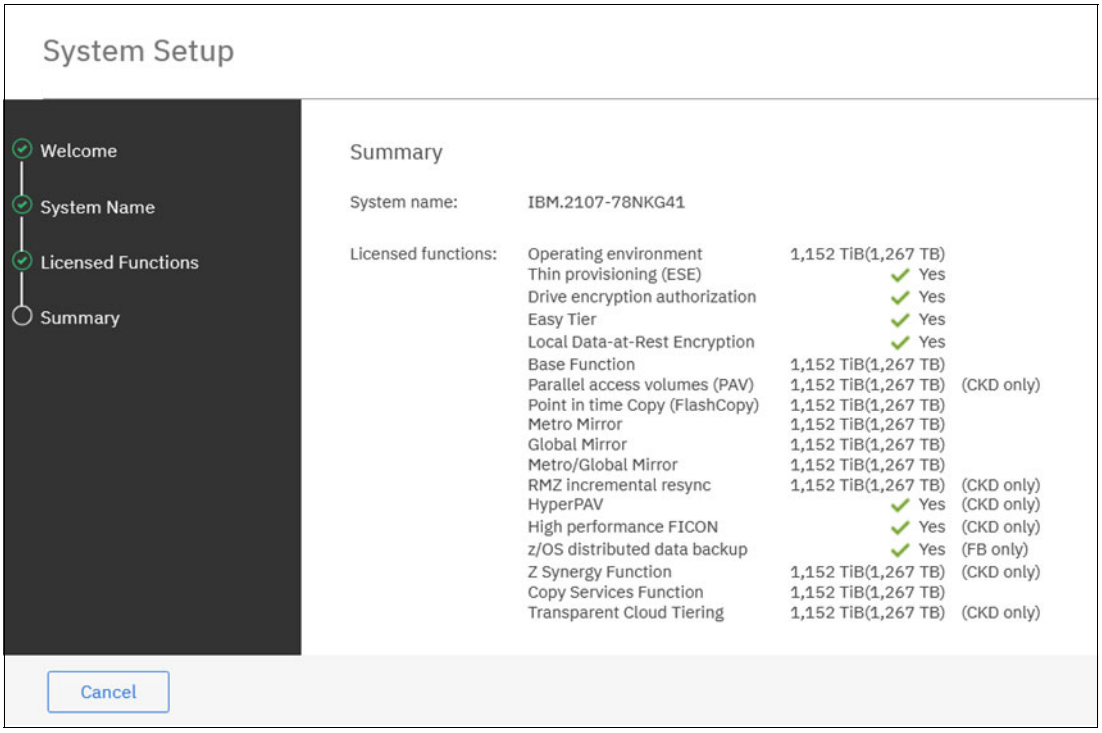


Figure 9-7 System Setup: Summary

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 on page 236.

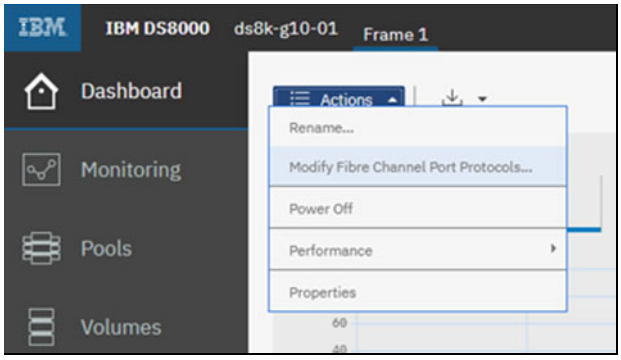


Figure 9-8 Quick way to modify the fiber adapter port protocols

- 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.

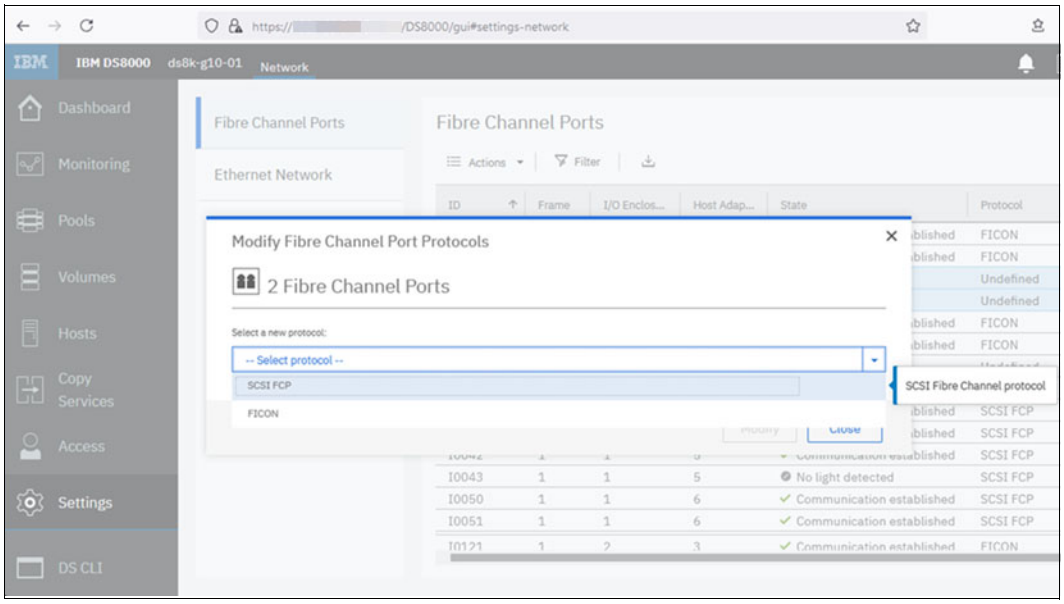


Figure 9-9 Configuring the Fibre Channel port topology

- You can also configure the port topology from the System view or from the System Health overview, as shown in Figure 9-10 on page 237.

You can also perform this configuration during logical configuration. For more information, see 9.7.4, “Creating FB host attachments” on page 279, and 9.14, “Monitoring system health” on page 305.

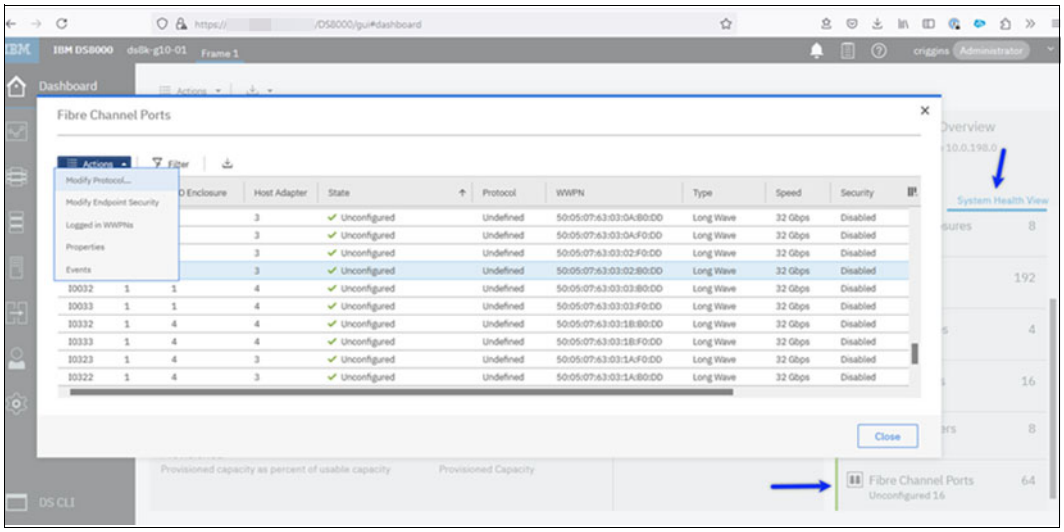


Figure 9-10 Modifying the FC port protocol from the System Health overview

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 DS8A00 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.

9.3.1 Overview of DS GUI

From the DS GUI, the administrator can manage the system, which includes performing the following actions for various activities:

- 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 Copy Services (CS) functions, such as FlashCopy and mirroring

Figure 9-11 on page 238 presents a high-level overview of the System window and all the objects that can be accessed from this window.

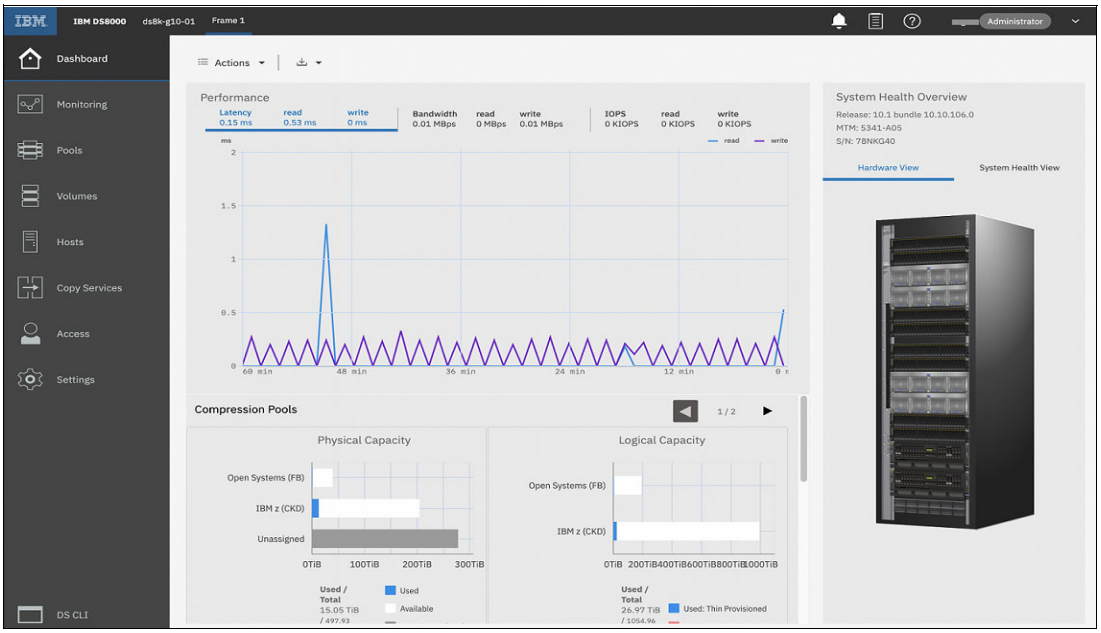


Figure 9-11 Initial System view

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 single view. The following items are included in this initial view:

- 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 DS8000 G10 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 on page 239.

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 IBM ESS 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

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 to 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.

The FC connectivity report 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 for which you should look into 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, such as code level, product, and machine serial number.

For more information about the DS8A00 help functions, see 9.3.2, "Storage Management help functions" on page 244.

– Hardware View:

A graphic of the DS8A00 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
- DS8A00 release
- Data that is specific to an IBM Power server

For more information about the Hardware view, see 9.14, "Monitoring system health" on page 305.

– 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 View

Clicking this tab opens the System Health View 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 (HAs), Device Adapters (DAs), and Fibre Channel 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. Below this area is the Provisioned Capacity summary of the system. To the right of this is the Capacity Savings for Open Systems showing the percentage and TiB saved via the use of Thin-Provisioning and Compression.

- 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. This view provides information about the pools within 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 lists any unassigned 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. Displays 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 allocated 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. Displays the hosts and host clusters defined in 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 logical path definitions.

Note: There are other, more powerful, Copy Services tools to review and manage Copy Services relationships such as IBM Copy Services Manager (CSM) and Geographically Dispersed Parallel Sysplex (GDPS).

- 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 user 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. This helps to ensure that the security of the system is more robust against internal attacks or mistakes. The following actions can be taken with 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 DS8000 G10 features and licensed functions” on page 195.

- 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) is also managed here.

- 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, date, and time manually 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 Copy Services, 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 choose 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 the embedded DS CLI is the most current one that is available for the microcode that is installed on the DS8A00.

– The DS CLI commands can be run 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.

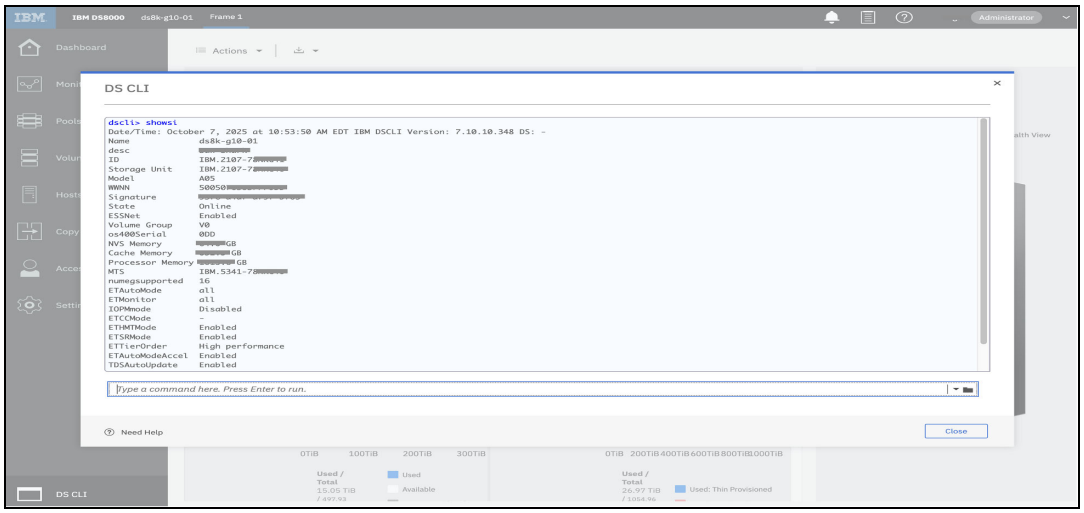


Figure 9-12 Embedded DS CLI

9.3.2 Storage Management help functions

The DS GUI provides access to comprehensive help functions. The help function aids you in using the GUI and provides in-depth details about the overall DS8A00 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 9-13 on page 245.

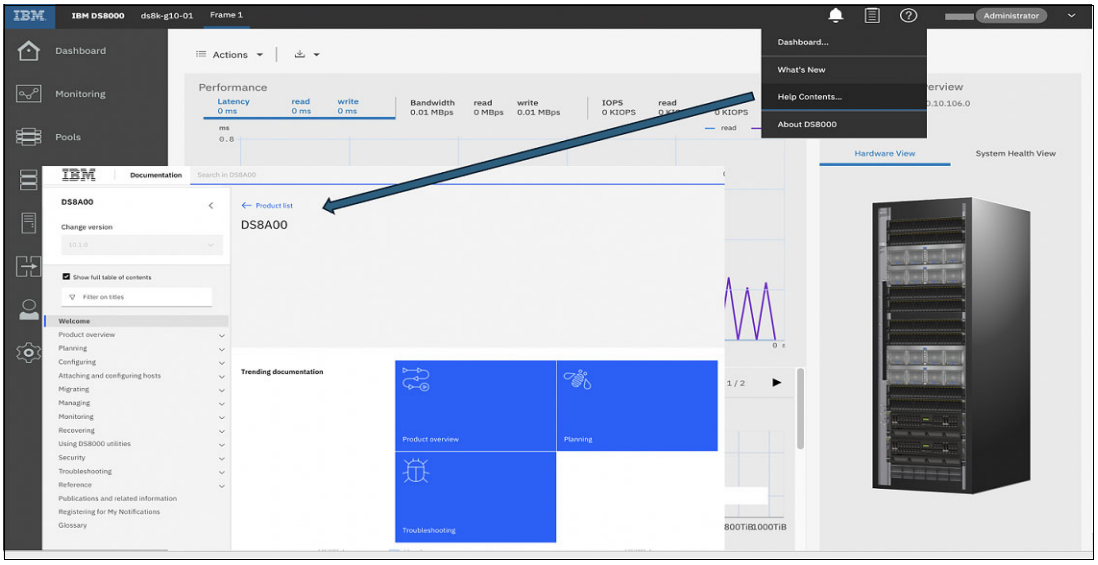


Figure 9-13 Storage Management Help System: IBM Documentation

In IBM Documentation, you can discover introductory information about the DS8A00 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 DS8000 documentation.

9.4 System configuration overview

After the DS8A00 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 DS8A00 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 9-14 on page 246.

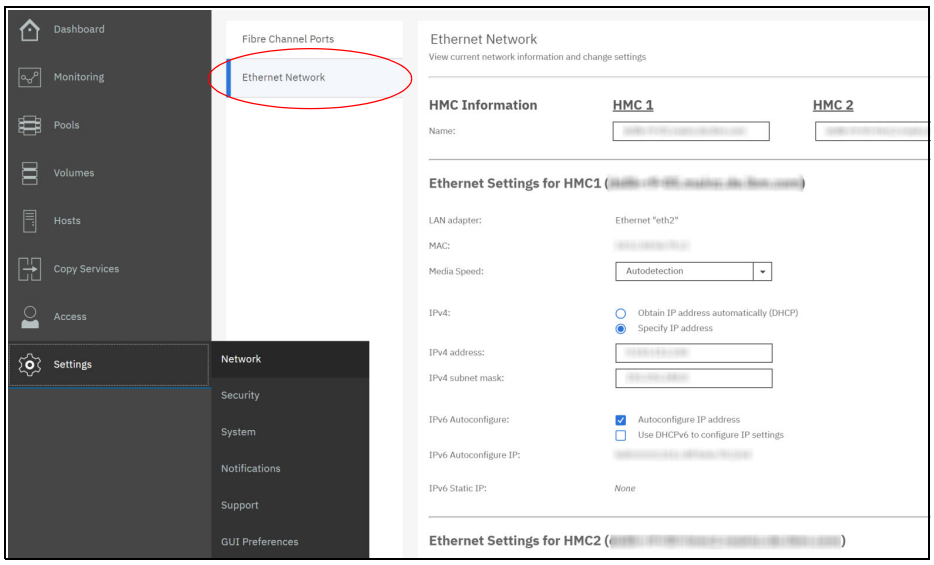


Figure 9-14 Ethernet Network settings

Fibre Channel ports

From the Fibre Channel Ports tab of the Network Settings window, which is shown in Figure 9-15 on page 247, you can configure the protocol that is used by the FC ports or view the properties of those ports. To view the Actions, select an FC port, click **Actions** → and select a function. You can also right-click on a specific port to see the same actions.

- Modify Protocol

  Set the protocol that is used for connecting to a host or another storage system.
- Modify Endpoint Security

  Configure IBM Fibre Channel Endpoint Security level: Disable, Enable, or Enforced.
- Refresh Security

  Port refresh causes the port to go offline for a few seconds and come 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 9.16, “Fibre Channel error rate statistics” on page 333.
- 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 329.
- 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 to your laptop or workstation that contains all the information for FC ports that are shown in the file, including their worldwide port names (WWPNs).

Figure 9-15 on page 247 shows the FC ports window with all available options that are listed.

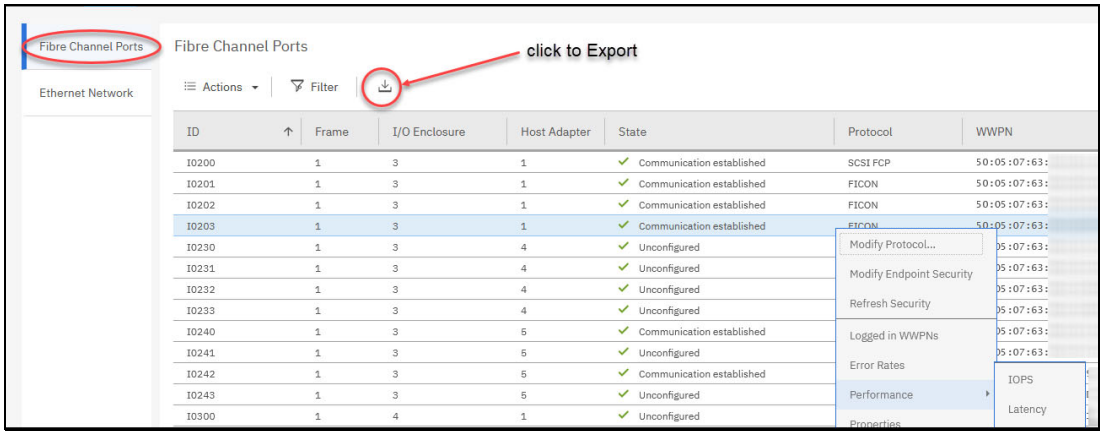


Figure 9-15 Fibre Channel Ports settings

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 DS8A00 offers the following encryption features:

- Data-at-Rest Encryption. Drive-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 DS8000 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 standards 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 DS8A00 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 9-16.

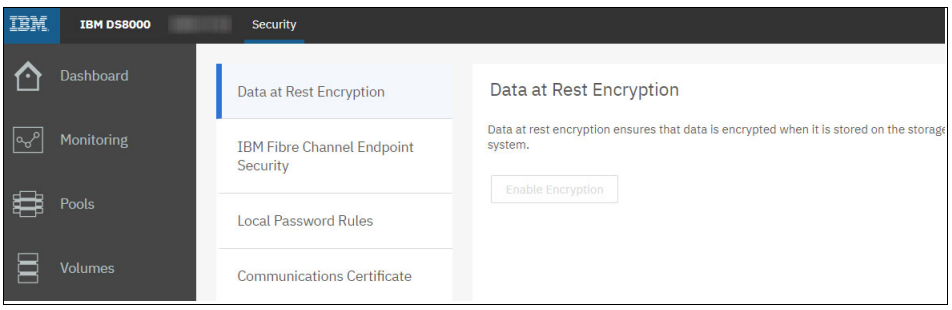


Figure 9-16 Enabling data-at-rest encryption

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 enabled and the encryption group is configured before starting any logical configuration. Once pools are created, encryption cannot be enabled or disabled.

If you use the Local Key Management feature (included in the base license), the DS8000 system manages the key group. Local Key Management can only be configured through the DS CLI. For more information, see *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security*, REDP-4500.

IBM Fibre Channel Endpoint Security

IBM Fibre Channel Endpoint Security protects data in flight between an IBM Z host and the DS8A00 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 DS8000 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 **Configure IBM Fibre Channel Endpoint Security**, as shown in Figure 9-17. This action starts a wizard that guides you through the process.

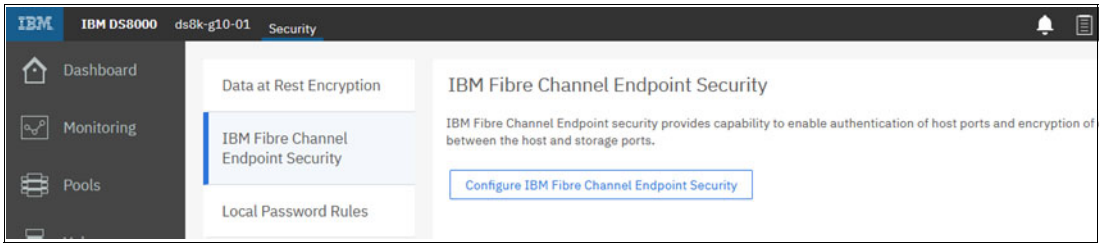


Figure 9-17 Configuring IBM Fibre Channel Endpoint Security

Local password rules

Many companies have their own security policies for passwords and want to implement the same policies for their DS8A00 storage system.

To implement password rules, complete these steps (see Figure 9-18 on page 249):

1. From the system window, click the **Settings** icon.
2. Click **Security** to open the Security window.
3. Click the **Local Password Rules** tab.

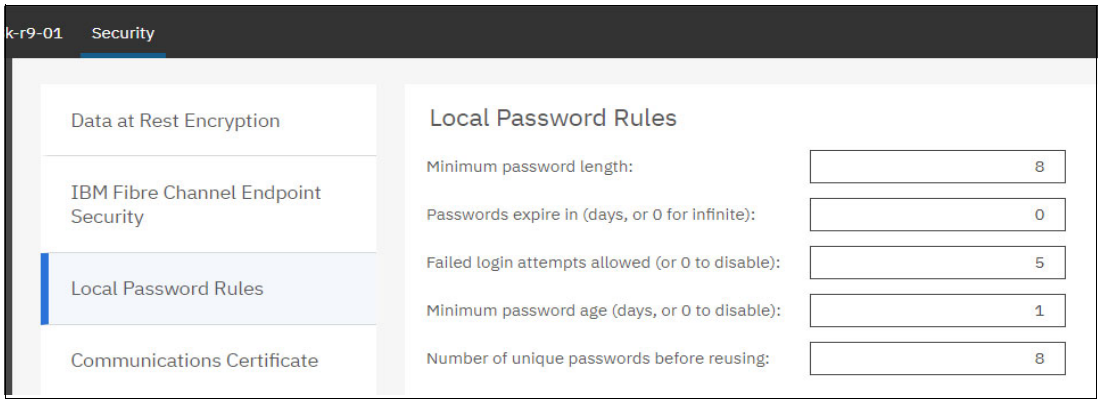


Figure 9-18 Local Password Rules

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 9-19.

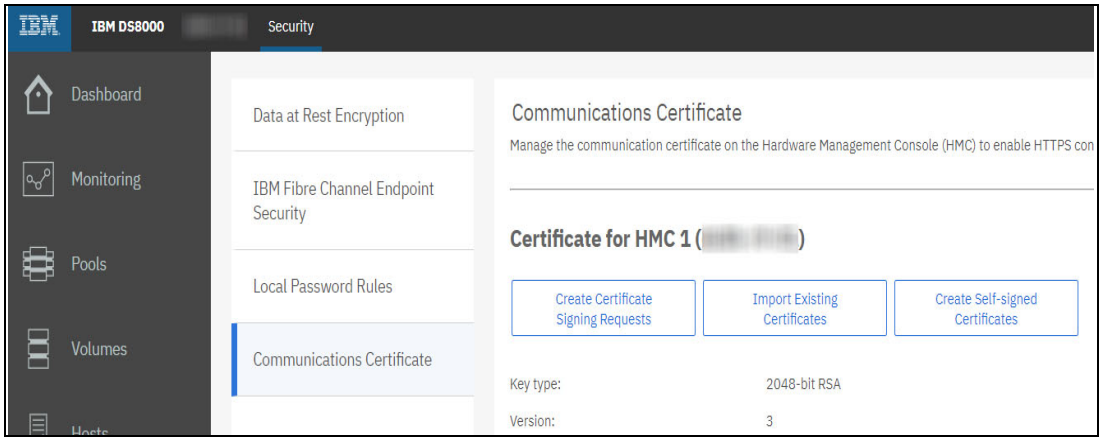


Figure 9-19 Communication Certificates

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 9-20 on page 250, 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.



Figure 9-20 Certificate signing request

After the CSR file is created, you can download that file for processing with your trusted CA.

Two other options are available:

- Importing existing certificates from your security group within your organization.
- Creating a self-signed certificate.

9.4.3 Security settings

Dual Control is a security enhancement feature for DS8000 storage systems that requires two separate authorized persons to approve and execute critical operations.

Key benefits are as follows:

- Prevents data loss from intentional or unintentional errors.
- Protects against privilege escalation and denial of service attacks.
- Ensures critical tasks require multiple authorized approvals.
- Maintains productivity for non-critical operations.
- Provides comprehensive audit trails for security compliance.

1. To enable Dual Control, select the **Access** icon from the DS GUI navigation menu on the left. Click **Dual Control**, then click **Enable Dual Control** icon, as shown in Figure 9-21 on page 251.

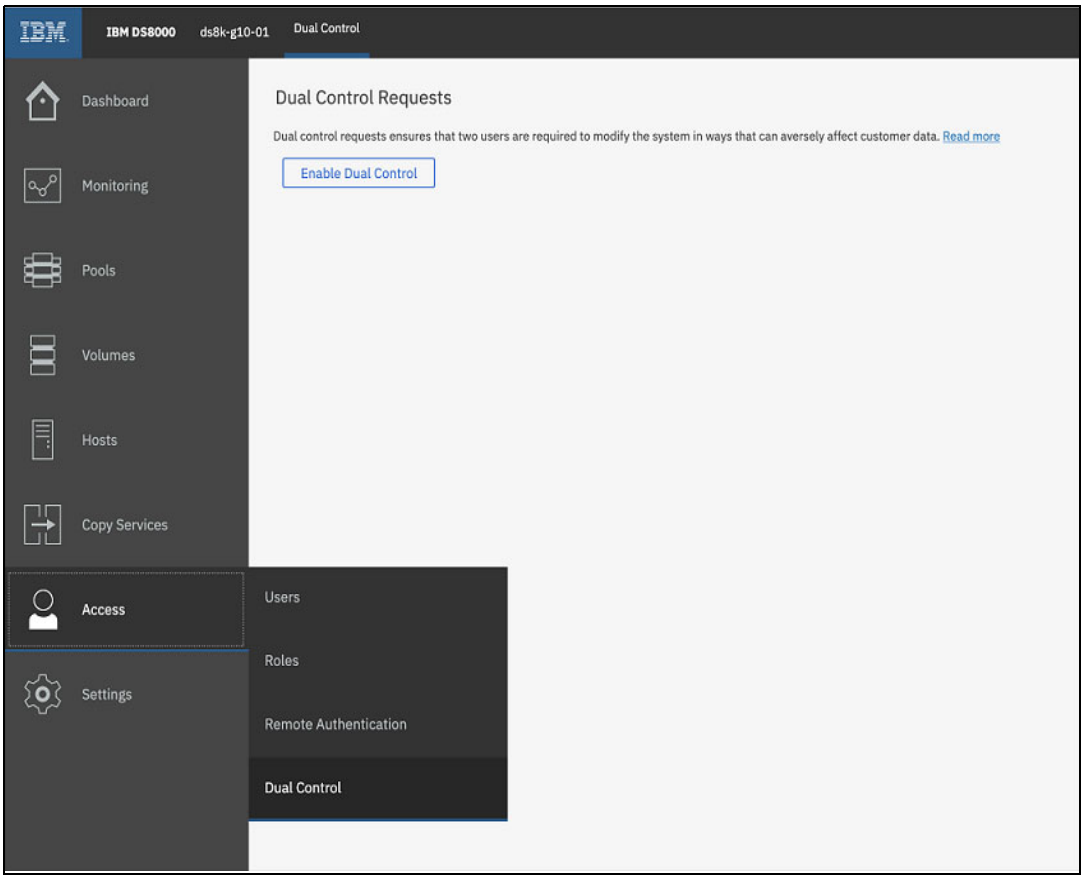


Figure 9-21 Dual Control Requests

2. Select the **Dual Control profile**. Click **Help** to see what tasks are protected under each profile, as shown in Figure 9-22.

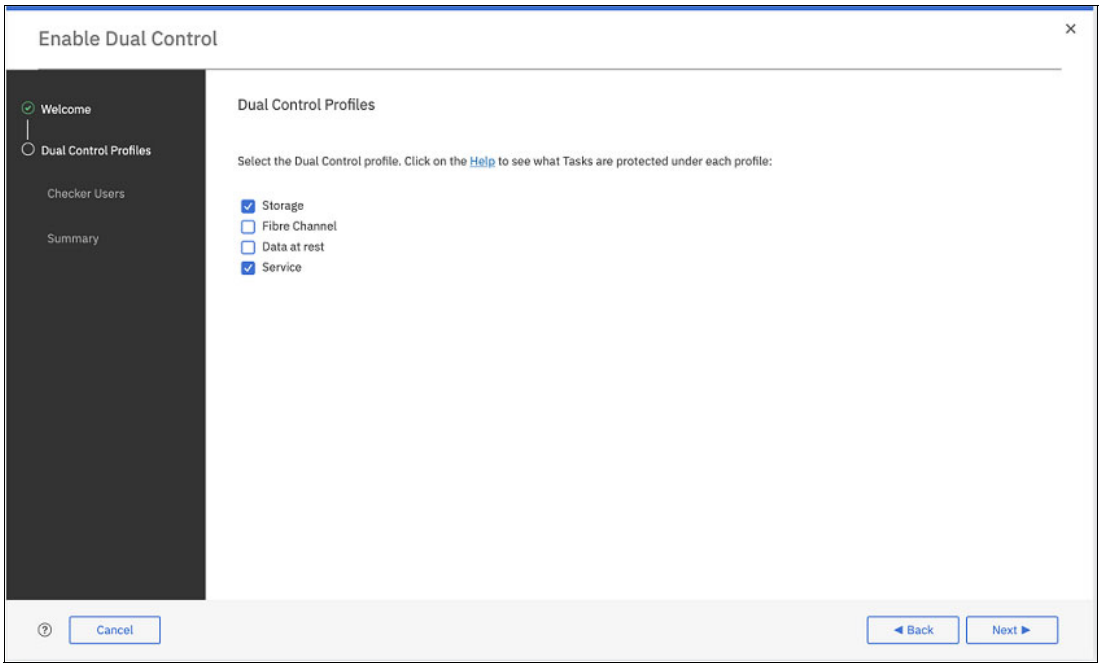


Figure 9-22 Dual Control Profiles

DS8A00 Dual Control – Protected Tasks are as follows:

- Permanent (Always Protected)
  - Enable/DisableDualControl
  - Create Administrator users
  - Modify user passwords (other users)
  - Grant/RemoveCheckerAuthority
  - Authentication policy changes
  - Remove DC protection from a user
- Service
  - Update HTTPS certificate
  - Modify network settings
  - System updates and services restarts
  - Service access modifications
  - Opening or closing ports
- Storage
  - Delete volumes and Safeguarded Copy capacity
  - Reinitialize volumes
  - Remove cloud servers
- DaR (Data at Rest)
- Data at Rest Encryption operations
- Key server management
- FC (Fibre Channel)
  - IBM Fibre Channel Endpoint Security operations
  - Modify Fiber Channel port configurations
  - Key server management

3. Select which user accounts will have the authority to approve or reject requests for protected commands. See Figure 9-23 on page 253.

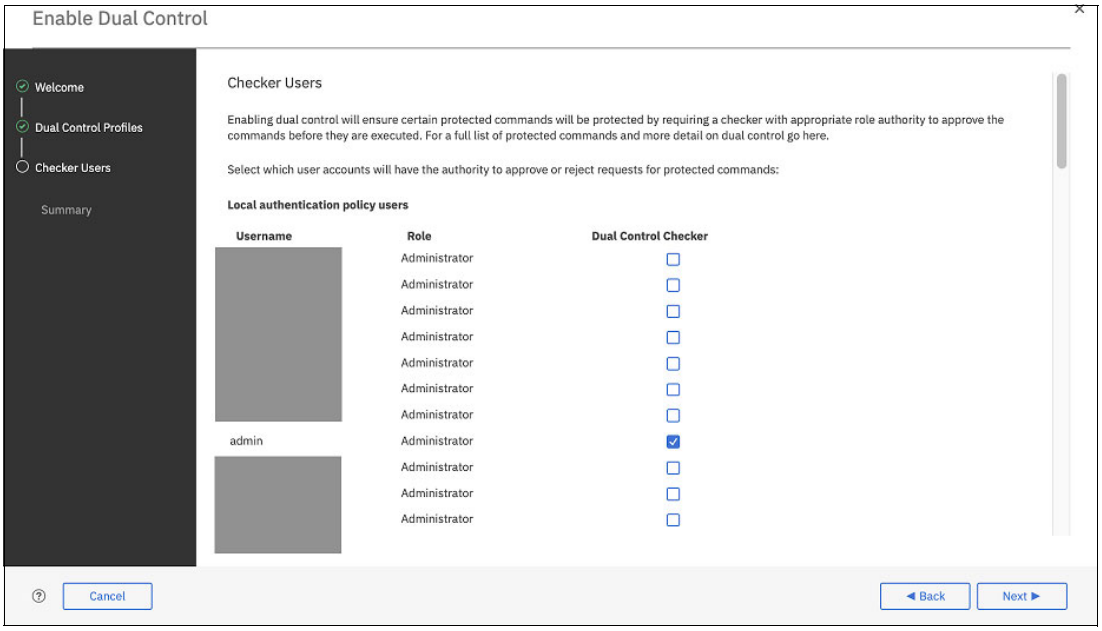


Figure 9-23 Enable Dual Control-Checker Users

4. Verify Dual Control settings and configurations before enable them, as shown in Figure 9-24.

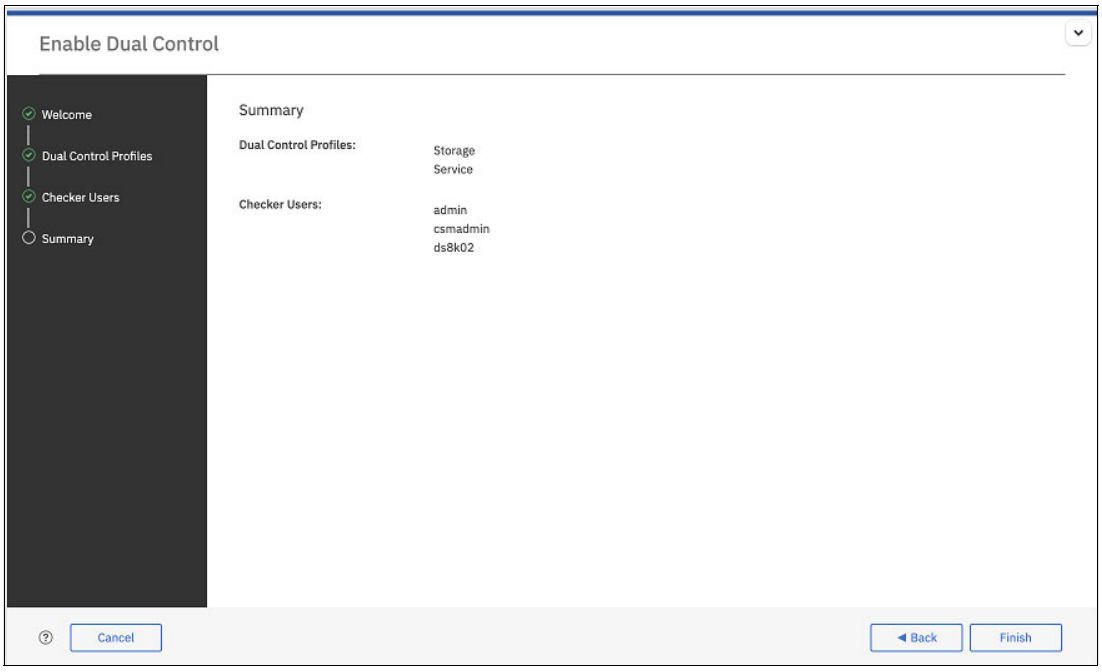


Figure 9-24 Finish Enable Dual Control Settings

5. If The new DSGUI Dual Control Panel shows all requests, as shown in Figure 9-25 on page 254 (If DS8A00 Dual Control is enabled.

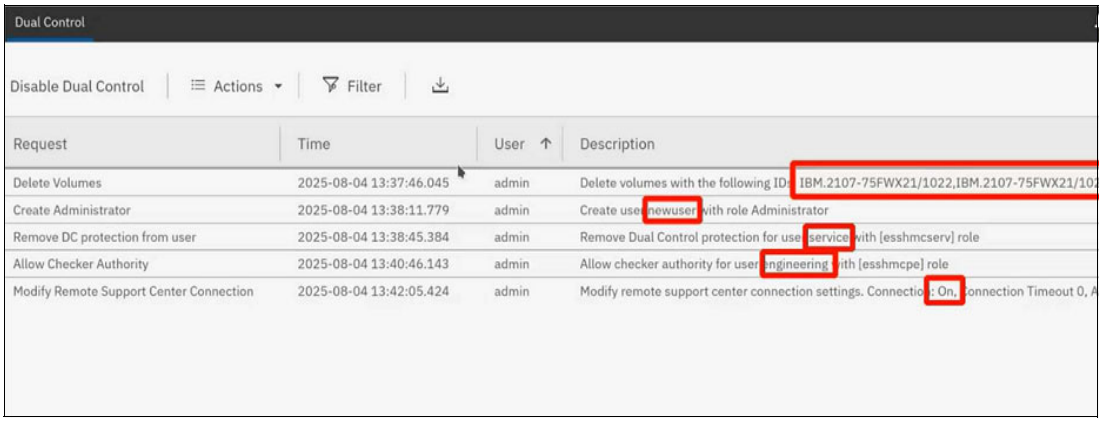


Figure 9-25 Dual Control requests

6. If a task is protected by Dual Control, the request displays the information shown in Figure 9-26 and Figure 9-27 on page 255 in the DS GUI.

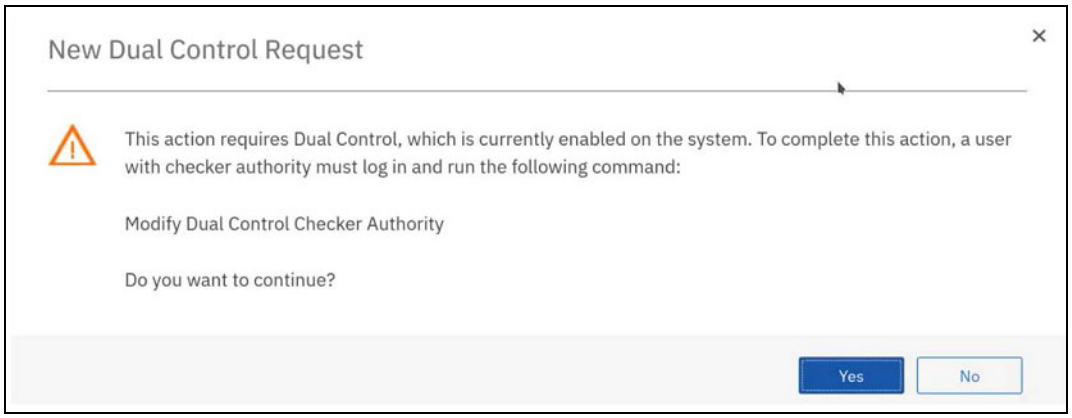


Figure 9-26 New Dual Control Request

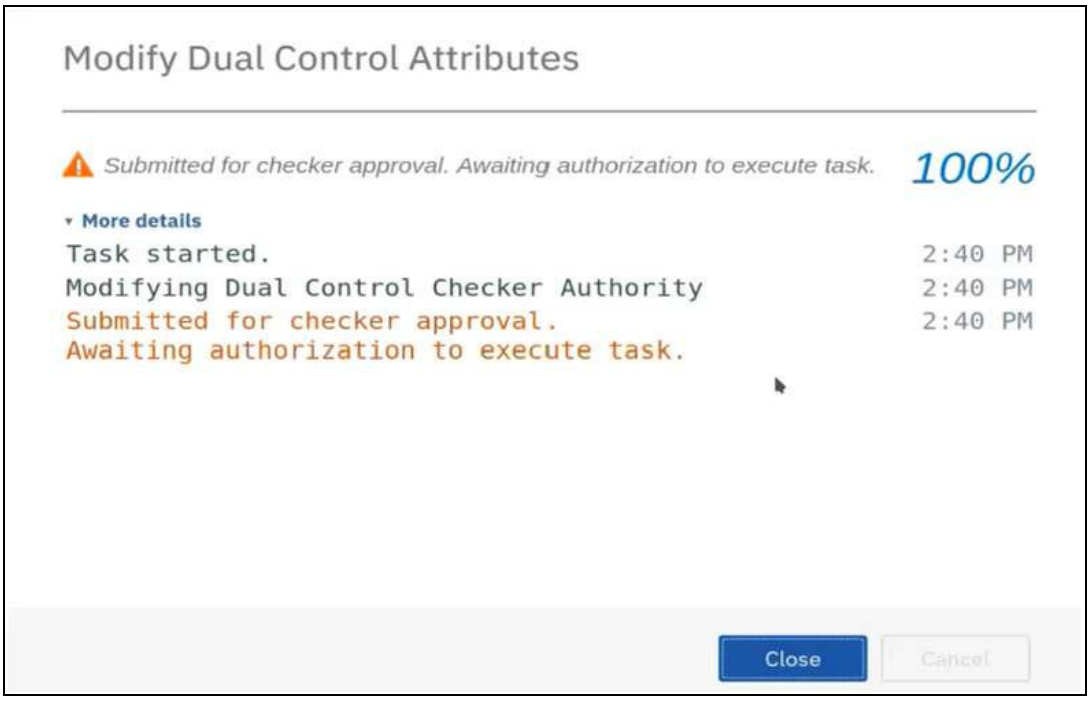


Figure 9-27 Modify Dual Control Attributes

7. The DS8A00 Dual Control checker (second authorized user who logs into the DS GUI to approve or reject the pending request) logs in to the DS GUI to approve requests, then clicks the **Task** icon and selects the **View** tab, as shown in Figure 9-28.

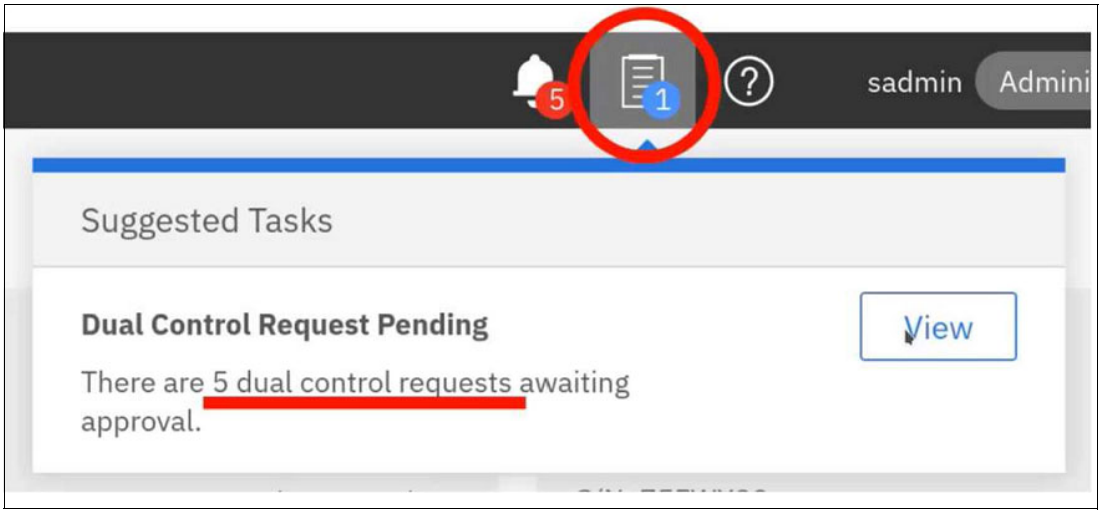


Figure 9-28 Suggested Tasks - Dual Control Requests

If the **Actions** tab is grayed out, it means either the checker cannot approve their own tasks or the user does not have the required authority. See Figure 9-29 on page 256 for reference.

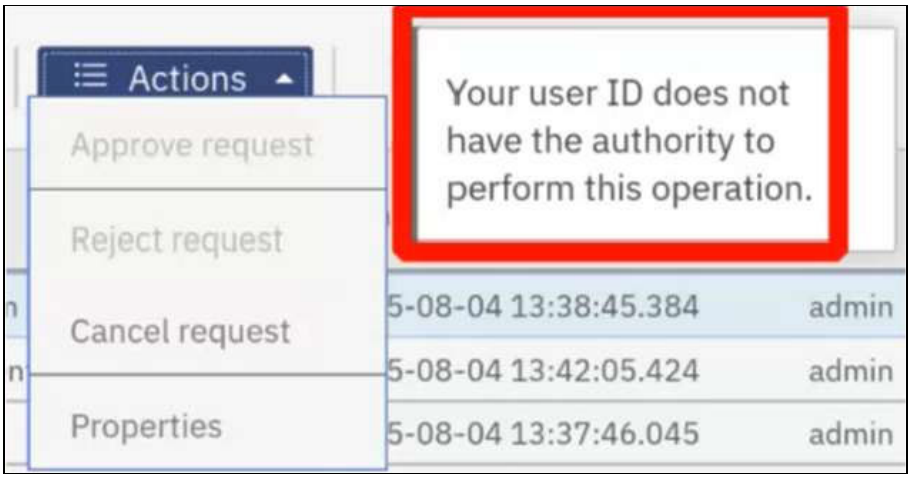


Figure 9-29 This user ID is not authorized to perform the requested operation

The user ID authorized to Approve or Deny the request seen in Figure 9-30 and in Figure 9-31 on page 257.

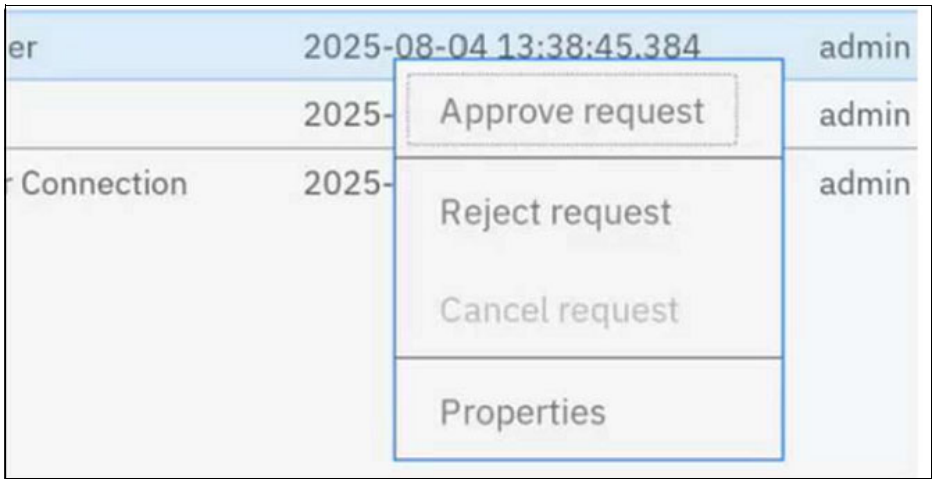


Figure 9-30 The user ID has the authority to perform this operation

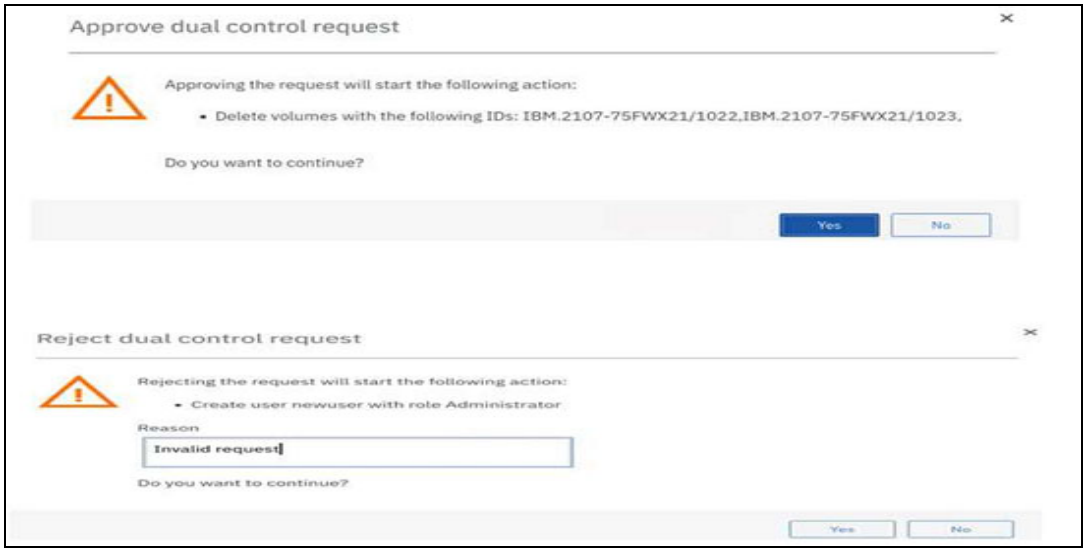


Figure 9-31 Approve or Reject dual control request

Once the Approve or Reject request is processed, it appears in the Events log, as shown in Figure 9-32 on page 257.

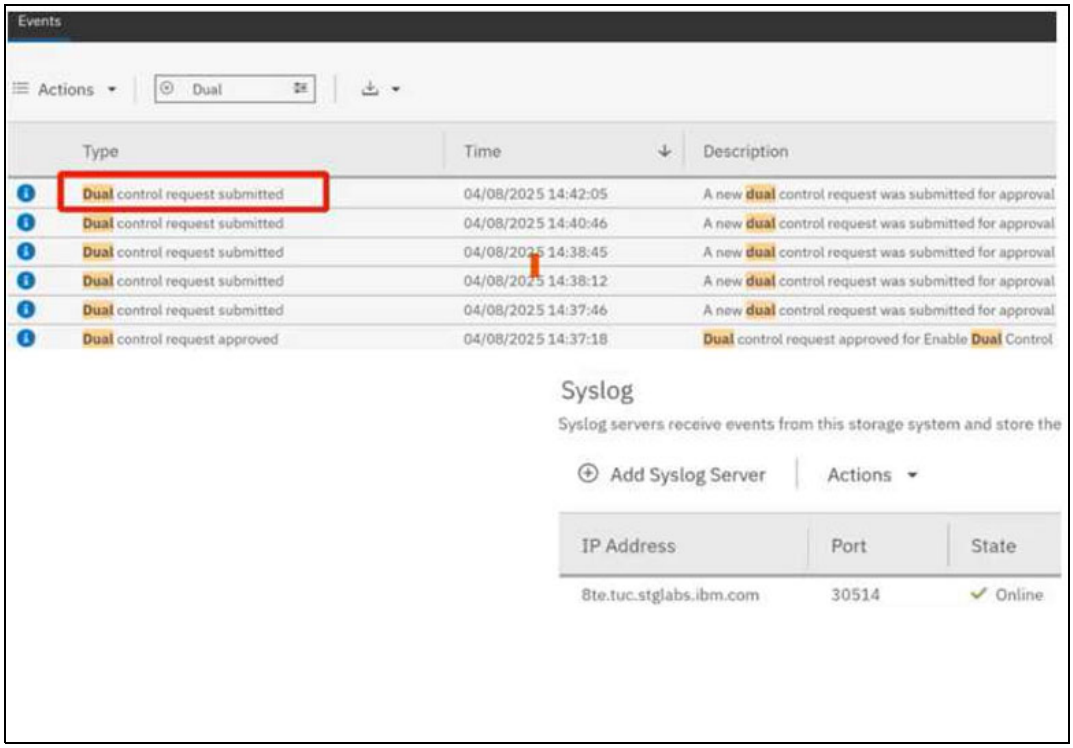


Figure 9-32 Dual Control request events

For more information, see IBM Documentation for Dual Control.

9.5 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 9-33.

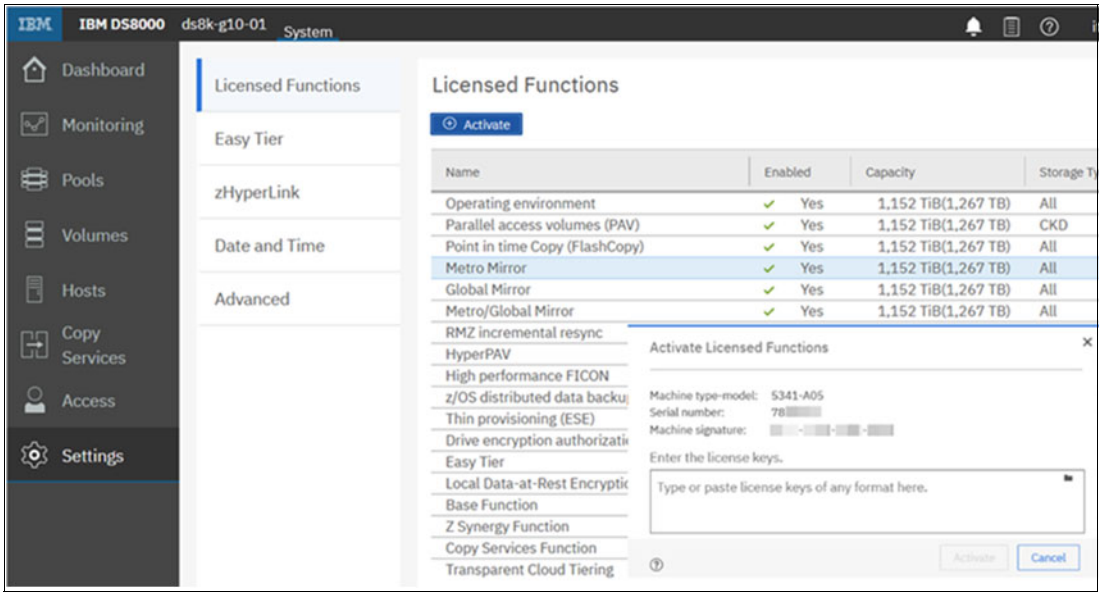


Figure 9-33 Licensed Functions settings

Easy Tier settings

To configure Easy Tier controls and advanced settings on your DS8A00 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 9-34.

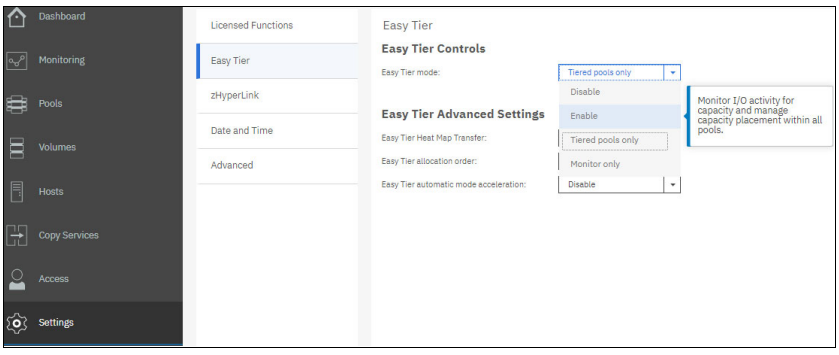


Figure 9-34 Easy Tier Settings

The following settings are available:

- Easy Tier mode. The available options are Disable, Enable, Tiered pools only, and Monitor only. When this setting is **Enabled**, 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*, REDP-4667.

zHyperLink

zHyperLink is a short-distance link (150 m) 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 9-35.

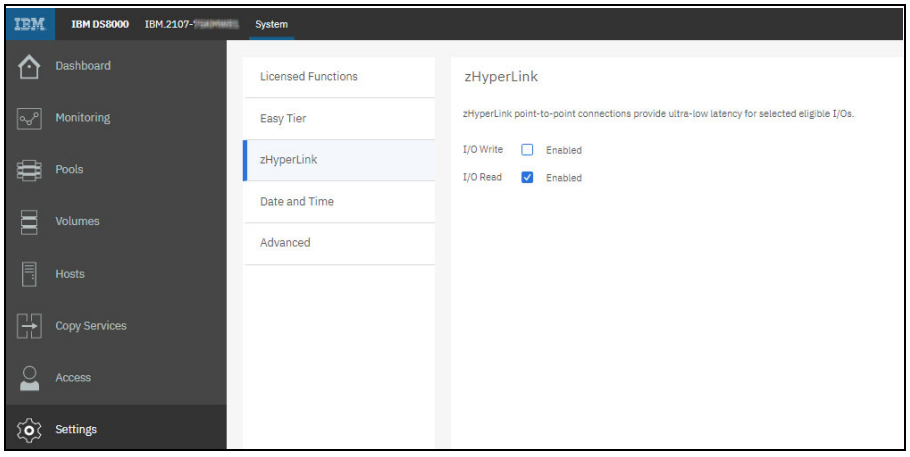


Figure 9-35 zHyperLink settings

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 9-36.

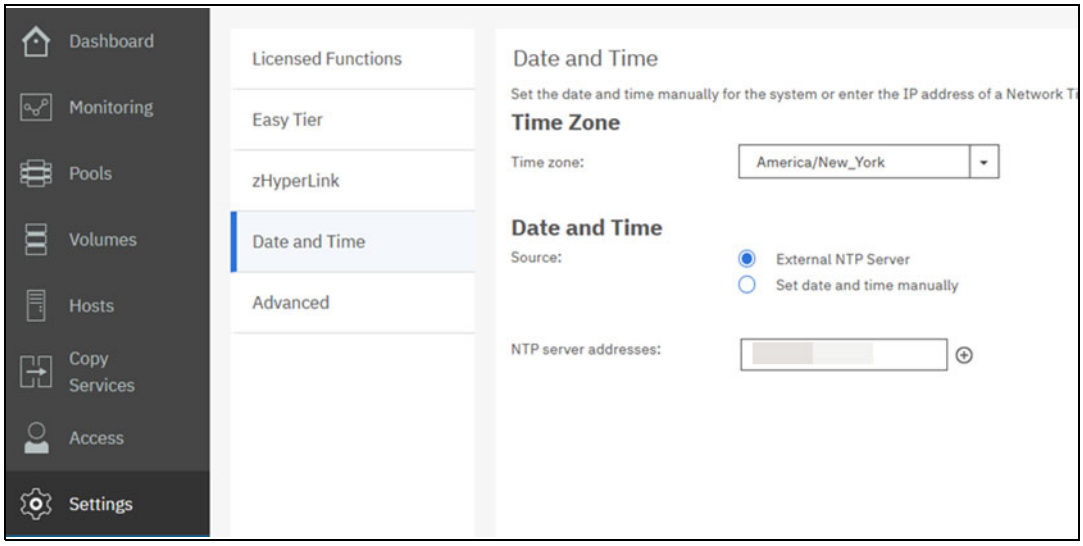


Figure 9-36 Date and Time settings

Advanced settings

From the **Advanced** tab of the System Settings window, you can configure system-wide options, such as power management, Copy Services, IBM Z features, and the DS Open application programming interface (API).

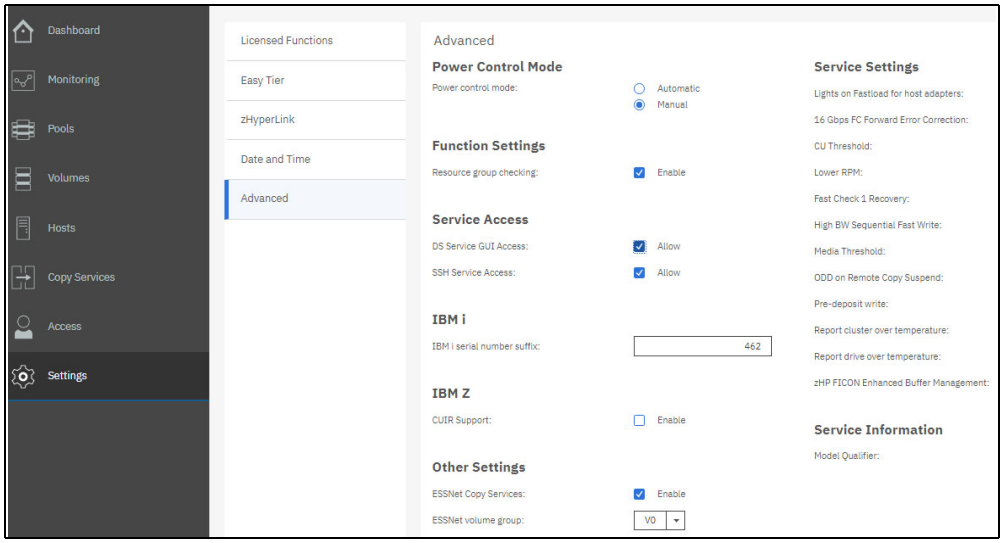


Figure 9-37 Advanced settings

The Advanced settings window is shown in Figure 9-37 and includes the following 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 9-37 on page 260). 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 Checking 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

  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

  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 Copy Services

    Enables the ESSNet user interface to manage Copy Services on the storage system.

  - ESSNet volume group

    Selects the ESSNet user interface to manage the volume group with Copy Services.

  - Host pre-check

    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 autonomy 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.

Enhanced Information Unit (IU) pacing uses the persistent pacing feature to change the operating pacing credit to 64. This benefits channel programs with CCW chains that contain more than 16 commands, or Db2 log writes in IBM Geographically Dispersed Parallel Sysplex (GDPS) multi-site workloads. For more information about this topic, see *IBM DS8000 and IBM Z Synergy*, REDP-5186.

- 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.

- Service Information

  Displays information that is related to the Model of the DS8A00.

9.5.1 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 9-38 on page 262.

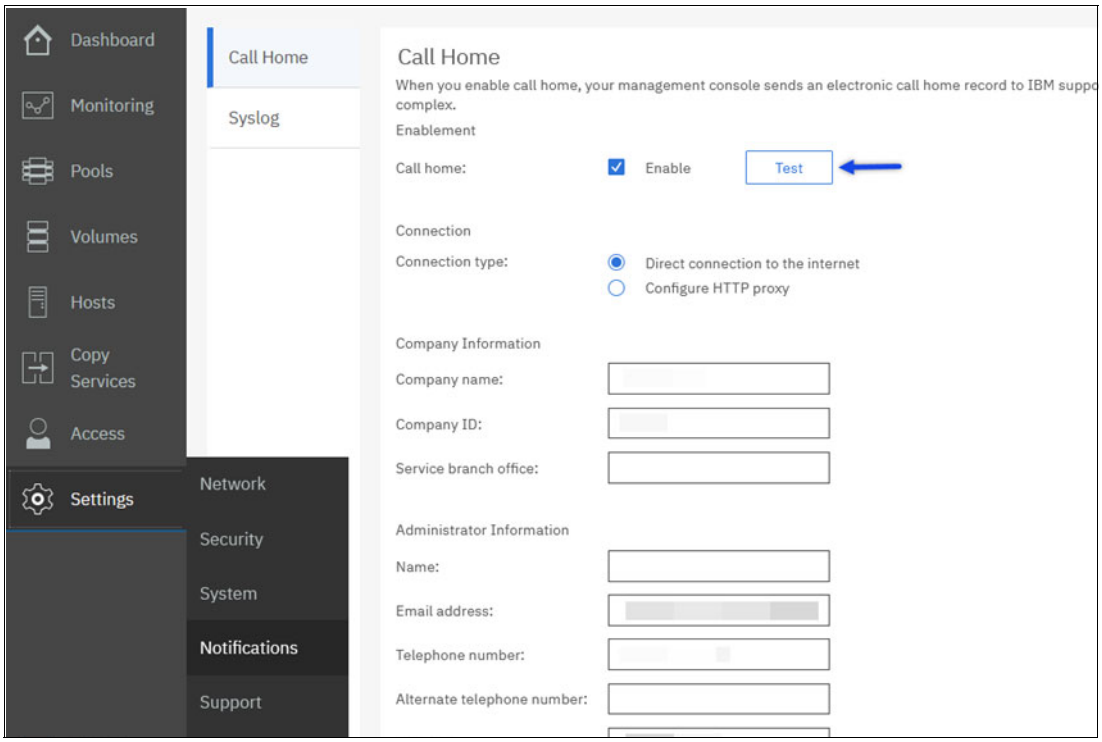


Figure 9-38 Notifications Settings window

Call Home

The DS8A00 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 9-38.

Important: All DS8000 storage systems must enable Domain Name Systems (DNS) on both HMCs or ensure both HMCs are able to access esupport.ibm.com through a firewall or proxy. Failure to do so results in Call Home not working.

Syslog

The Syslog window displays the Syslog servers that are configured to receive logs from the DS8A00 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 the 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 DS8A00 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 9-39.

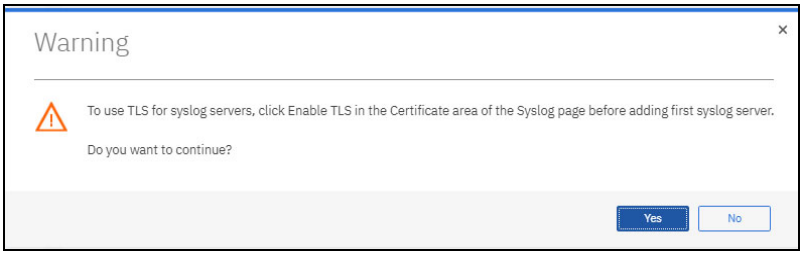


Figure 9-39 Warning to enable TLS

Note: A DS8A00 server must use TLS for its communications with the Syslog server. To configure TLS, the customer must generate their own trusted certificate for the DS8A00 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 9-40.

For more information about the setup of the SYSLOG server with TLS, see Encrypting Syslog Traffic with TLS (Secure Sockets Layer) (SSL).

This 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 9-40 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 9-40 on page 264.

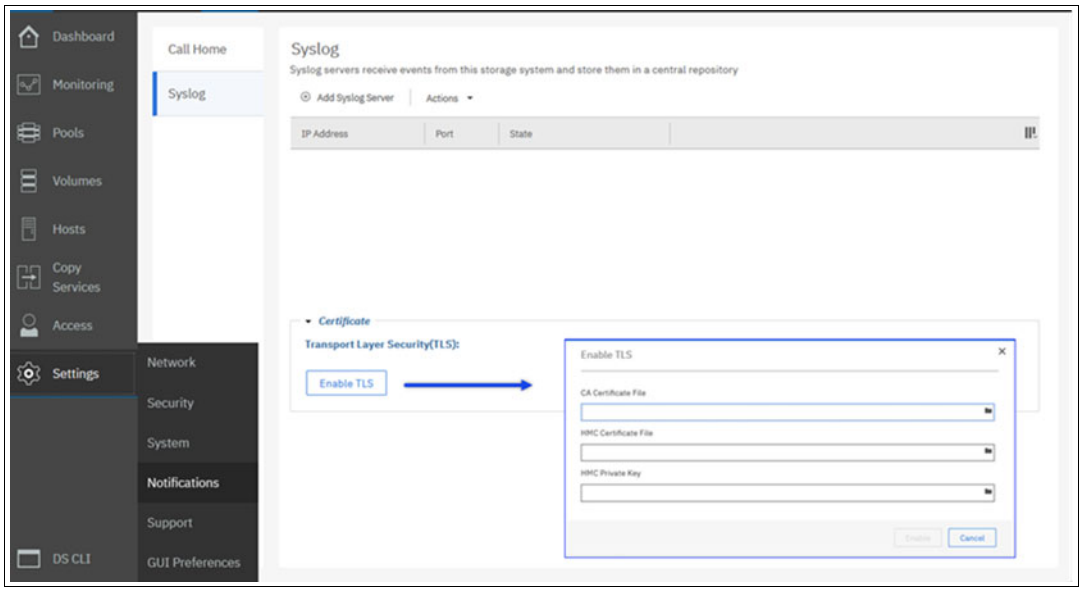


Figure 9-40 Enable TLS window

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 9-41, 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.

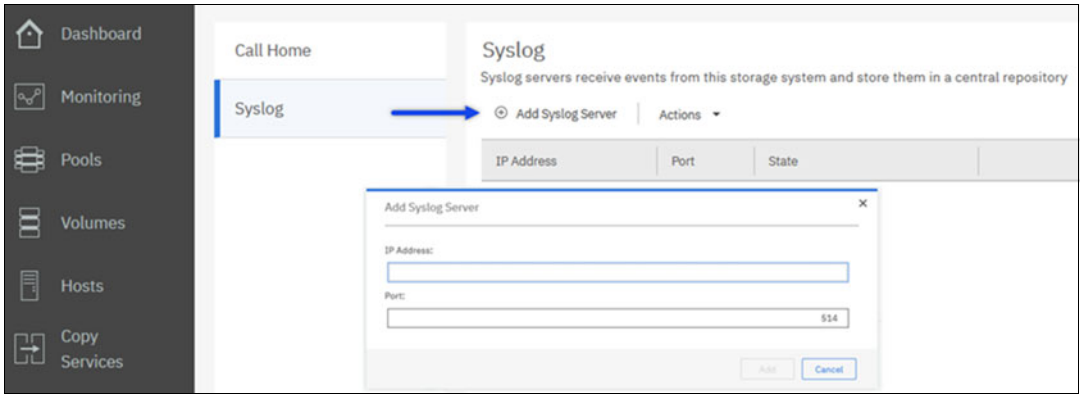


Figure 9-41 Syslog settings window

8. After the required Syslog servers are created, you can **Modify**, **Test**, **Activate**, **Deactivate**, and **Remove** a selected Syslog server, as shown in Figure 9-42 on page 265.

Note: To enable TLS, all existing Syslog servers must be deleted first. Then, you can enable TLS and create the Syslog servers.

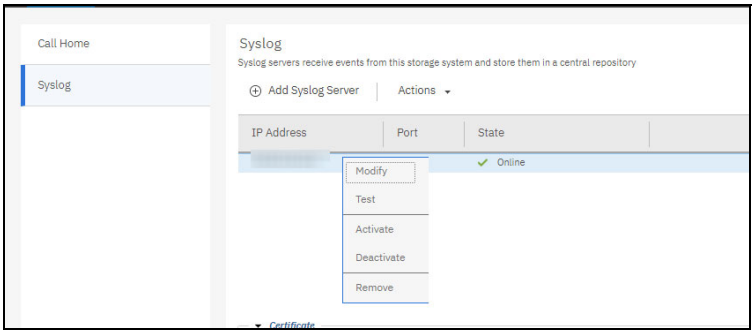


Figure 9-42 Modify Syslog configuration

9.5.2 Support settings

Use this section to configure various support settings.

Update System

You can perform a System Health Check on the storage system. In addition, you can view and manage the microcode levels that control all storage system functions.

IBM Remote Support Center

On the Support settings window, you can configure service access to the HMC so that RSC can access the HMC for problem determination, as shown in Figure 9-43. Note this figure displays the full configuration as if “Show full configuration...” was selected.

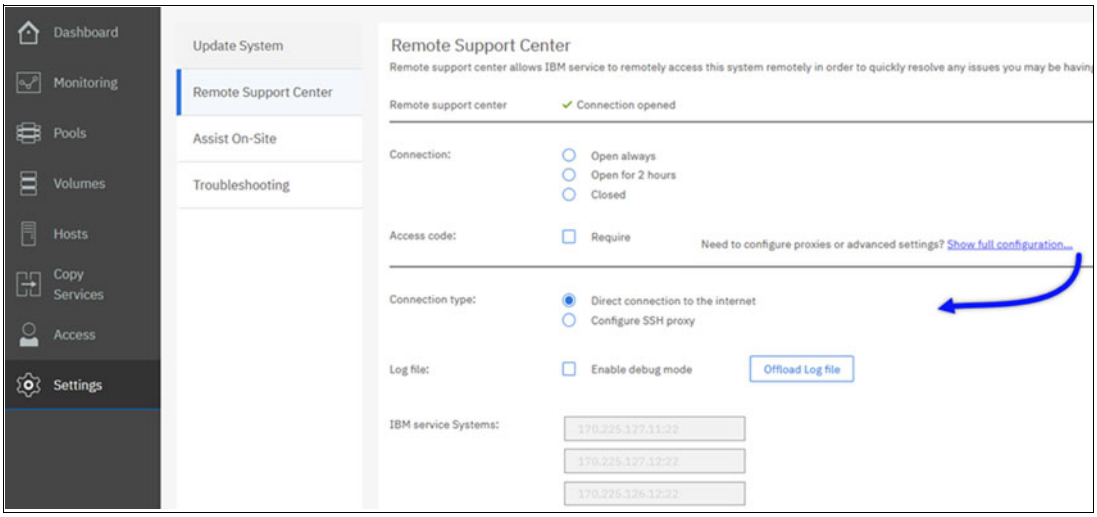


Figure 9-43 IBM Remote Support Center settings and full configuration

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**. From there, you can select the connection type, define proxy settings, and offload log files. The **Offload Logfile** option downloads the RSC log file, which contains a record of events and actions for the storage system. To generate detailed logs, select **Enable Debug Mode**. The **IBM Service Systems** section lists the IP addresses of the RSC servers; ensure your firewall is configured 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 9-44 on page 266.

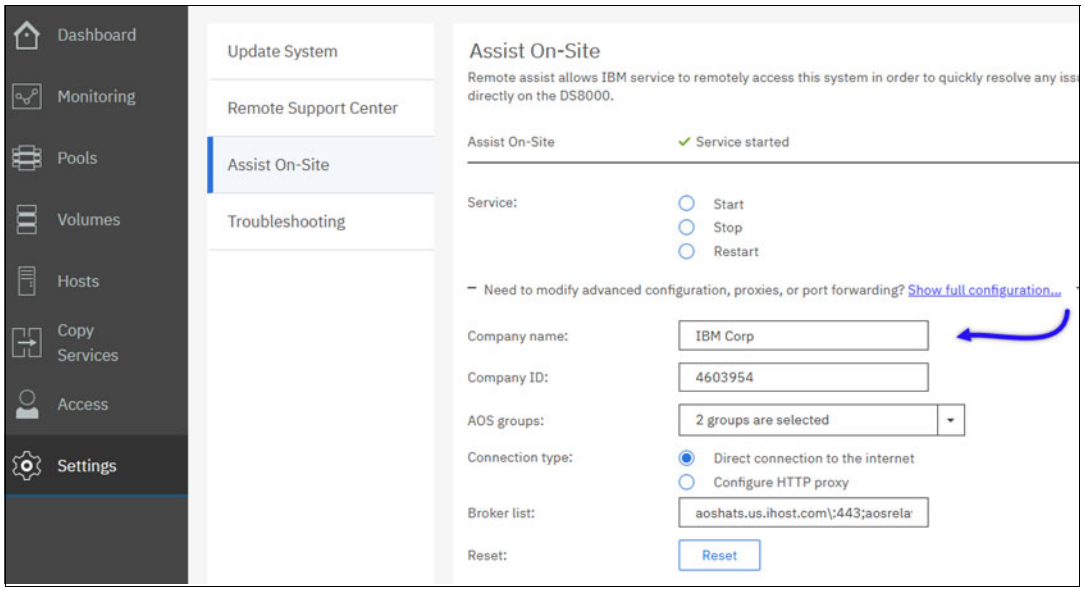


Figure 9-44 AOS settings and full configuration

To configure AOS, click **Show Full Configuration** and enter the required settings, as shown in Figure 9-44.

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 9-45 shows the troubleshooting options available.

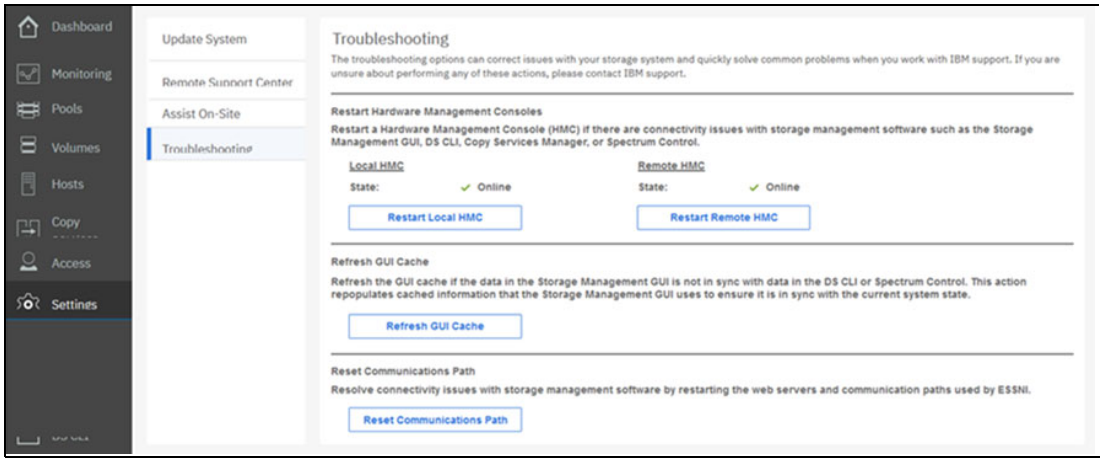


Figure 9-45 Troubleshooting options

GUI Preferences

Use the GUI Preferences tab that is shown in Figure 9-46 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.

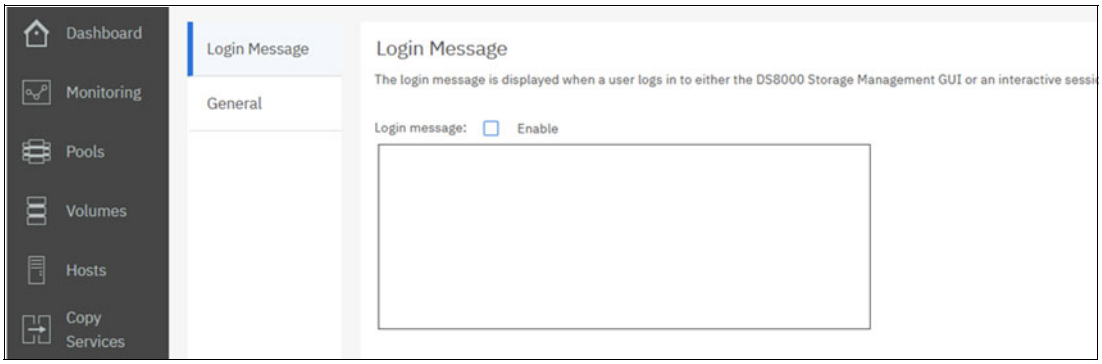


Figure 9-46 Login message

9.6 Logical configuration overview

The logical configuration of the DS8A00 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.

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

When you plan your configuration with the DS8A00, 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 9.8, “Logical configuration for Count Key Data volumes” on page 289.

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.

With the DS8A00 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 or 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 DS8A00 Storage Management Command-line Interface” on page 335.

Note: Deleting a pool with volumes is available via the GUI. A warning is displayed, and the user must enter a code that is presented by the DS8A00 to confirm the delete. A “force deletion” option is also available. For more information, see Figure 9-96 on page 302.

9.7 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

9.7.1 Configuration flow

The following steps 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.

9.7.2 Creating FB pools for open systems hosts

For best performance and a balanced workload, create at least 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 273.

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 9-47.

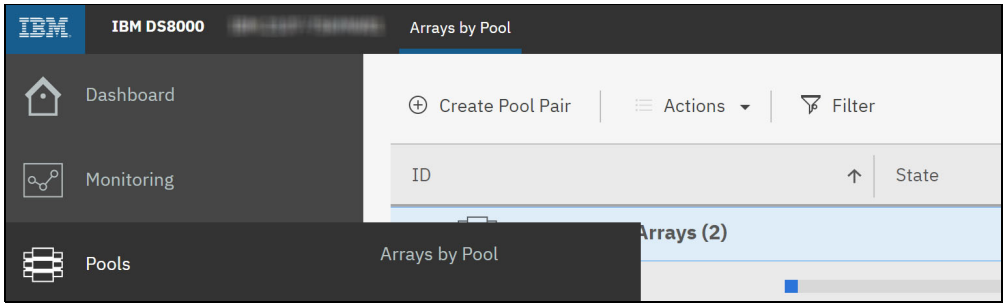


Figure 9-47 Arrays by Pool

2. Click the **Create Pool Pair** tab, as shown in Figure 9-48. The Create Pool Pair window opens.

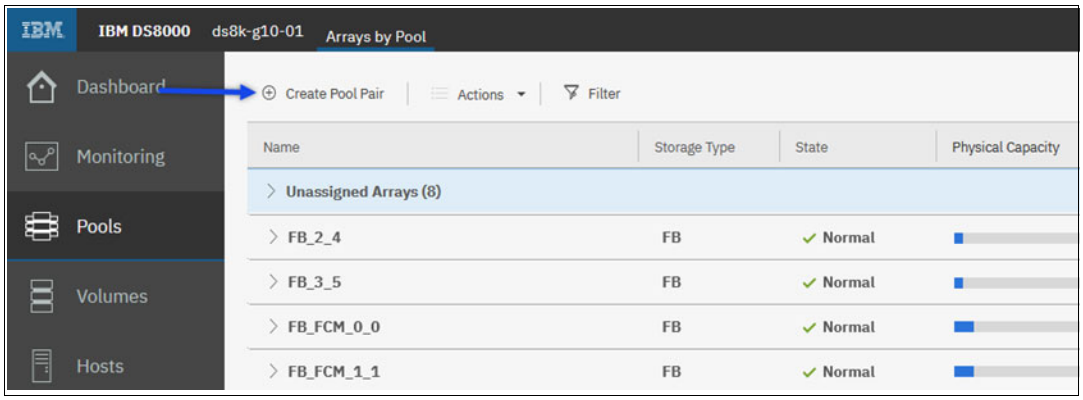


Figure 9-48 Create Pool Pair window

Note: By default, arrays within the DS8A00 are created using the virtualization technology of Redundant Array of Independent Disks (RAID). When an array is assigned to a pool, the choice of RAID 6 or RAID 10 is available.

3. Specify the pool pair parameters, as shown in Figure 9-49 on page 271:
   - 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.
4. Select whether the pair will be Compressed. Compression is important:
   - If no compression is selected, the industry-standard NVMe disks are used. Standard, full-provisioned, volumes can be created. No arrays are automatically assigned.
   - If compression is selected, FCMs are used. Only ESE volumes can be created. Arrays are automatically assigned during pool pair creation. See Figure 9-78 on page 290.

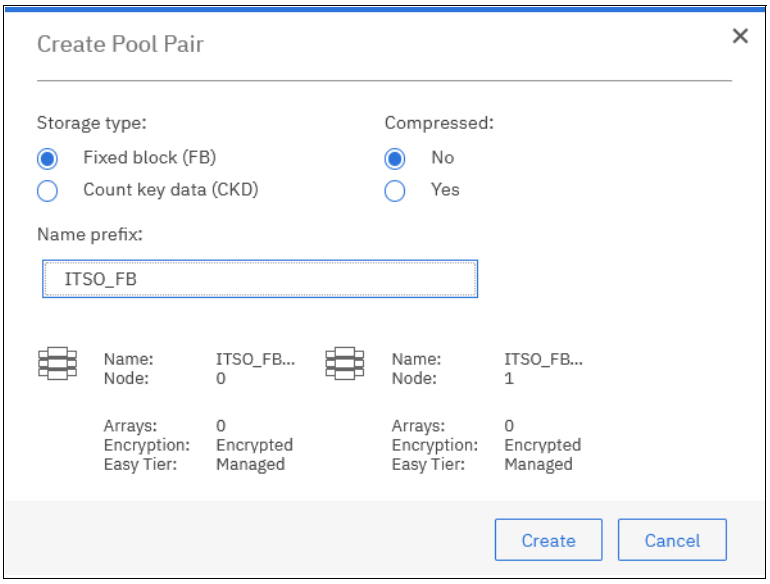


Figure 9-49 Creating an FB pool pair

5. When the pool pair parameters are correctly specified, click **Create** to proceed. Figure 9-50 shows a pool pair that is created.

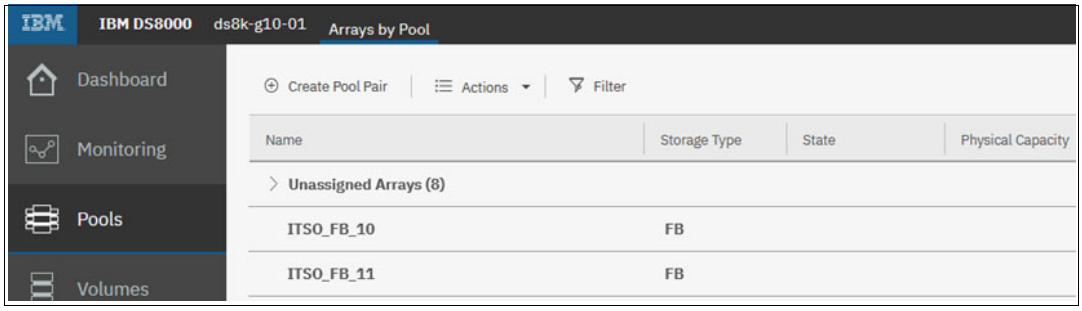


Figure 9-50 Pool pair that is created with assigned arrays

6. Assign arrays to the pool pair. Click **Unassigned Arrays** to display them, Right-click the array to be assigned and then **Assign** as shown in Figure 9-51 on page 271.

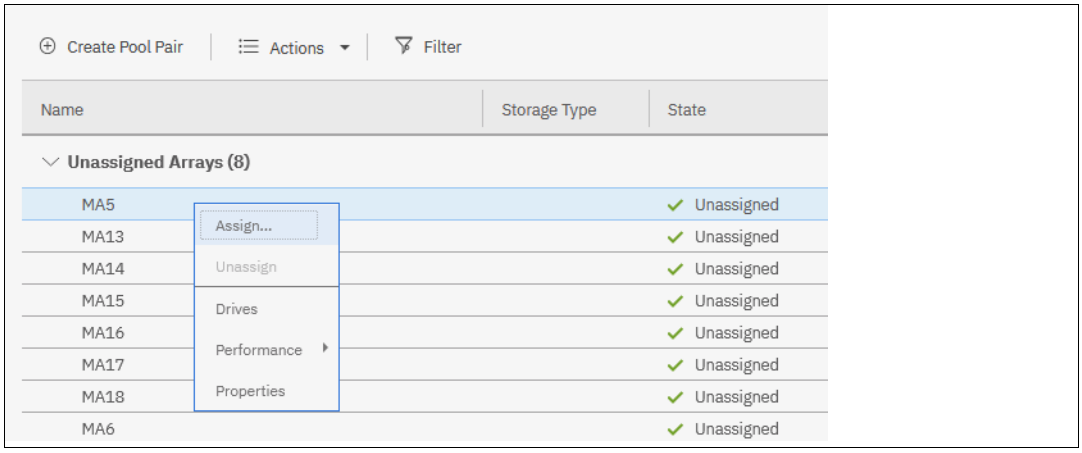


Figure 9-51 Unassigned Arrays and assigning them

7. The Assign Array window opens. Select the Pool from the drop-down menu that you want to assign the array. Choose the RAID type: 6 or 10. Select the Extent size: 1 GiB or 16 MiB. Click **Assign** to complete the process. Figure 9-52 is an example of assigning an array and Figure 9-53 on page 272 shows the pool pair with assigned arrays.

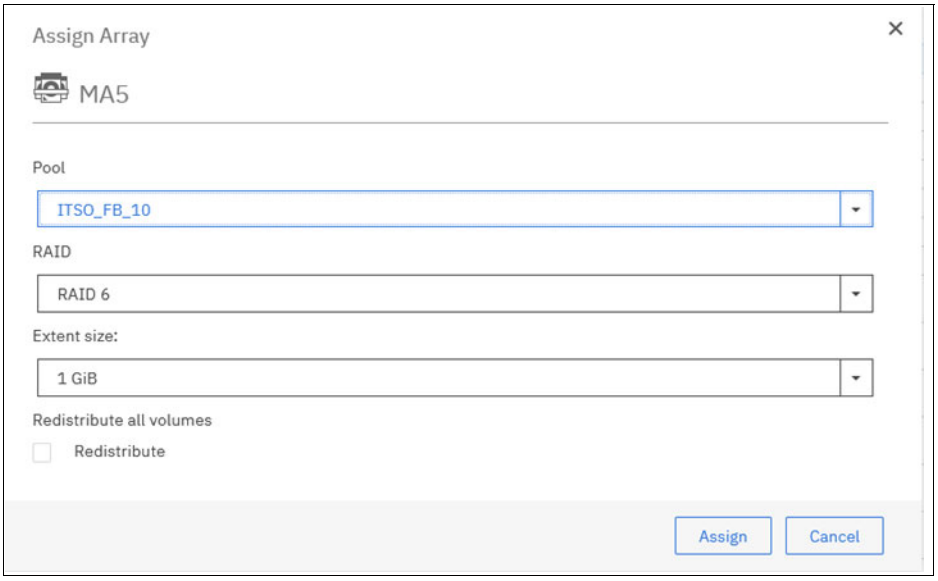


Figure 9-52 Assigning an Array to a Pool

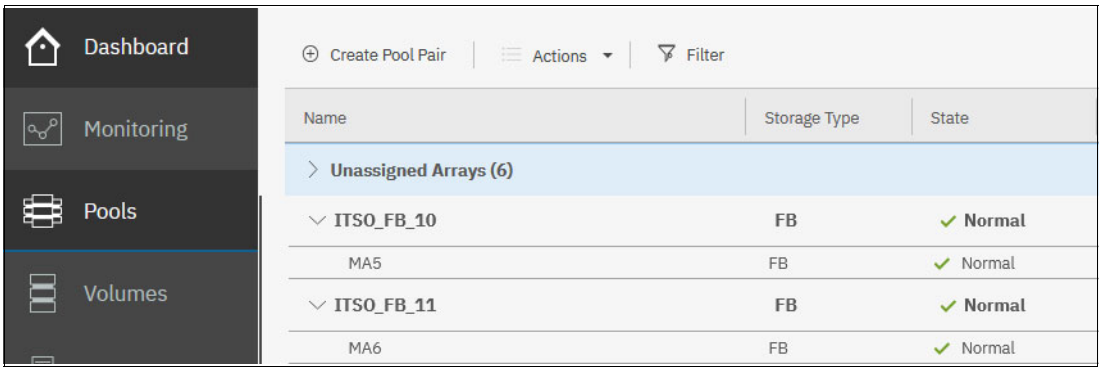


Figure 9-53 Pool pair with assigned arrays

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 right-click **Assign** or **Reassign**, as shown in Figure 9-54 on page 273. This action opens the Assign Array window.

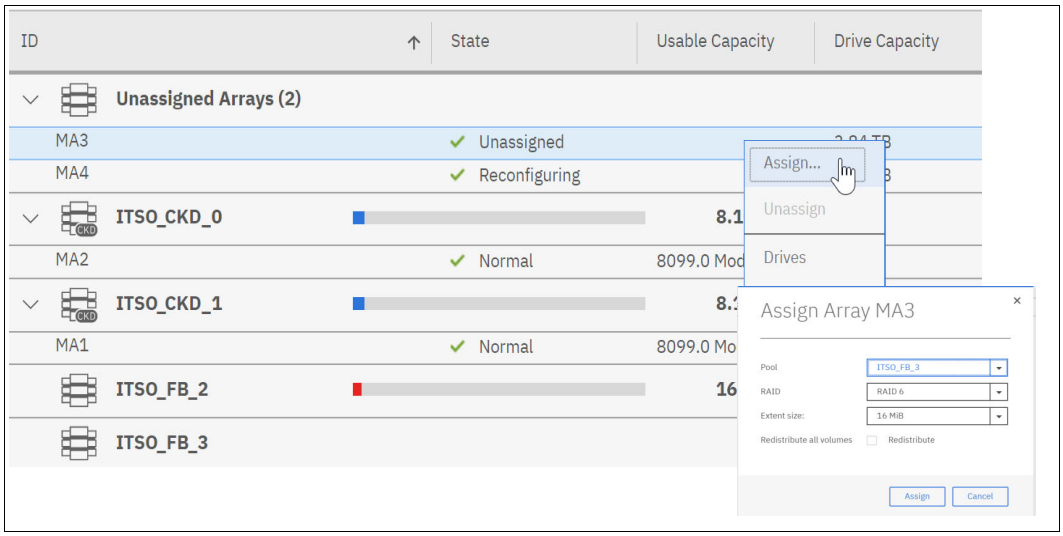


Figure 9-54 Manually assigning an array to an existing pool

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*, 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 9-55. By default, arrays are not automatically assigned to pool pairs.

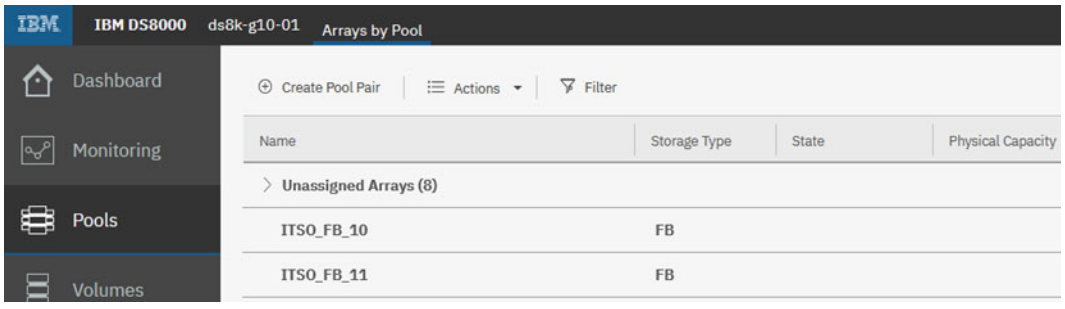


Figure 9-55 Creating an empty pool pair with no assigned arrays

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 9-56 on page 274.

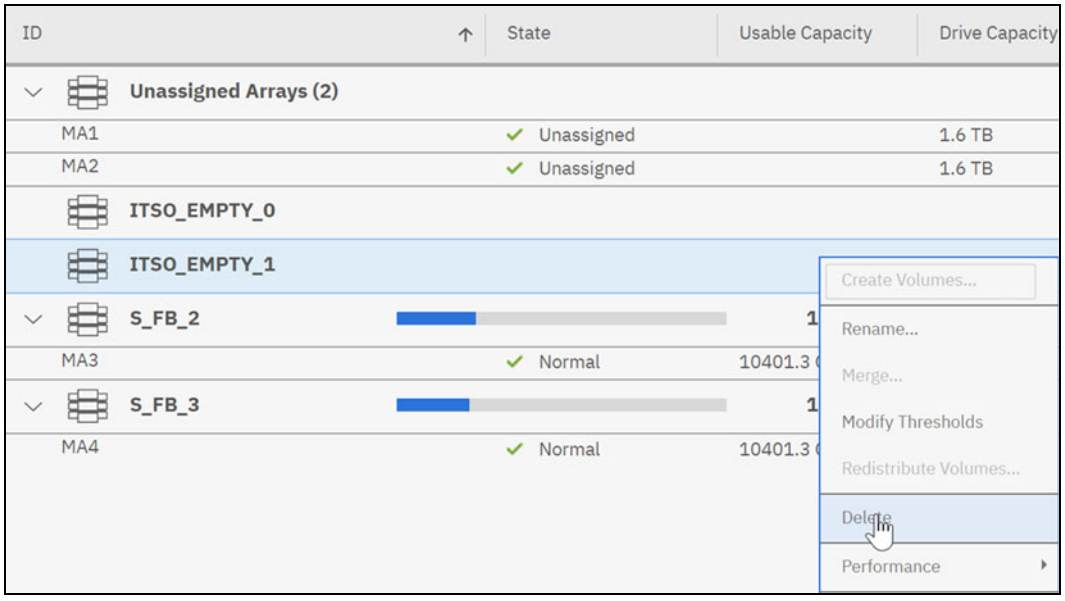


Figure 9-56 Deletion of a single pool of a pair

4. Assign one or more arrays to the single pool, as shown in Figure 9-57.

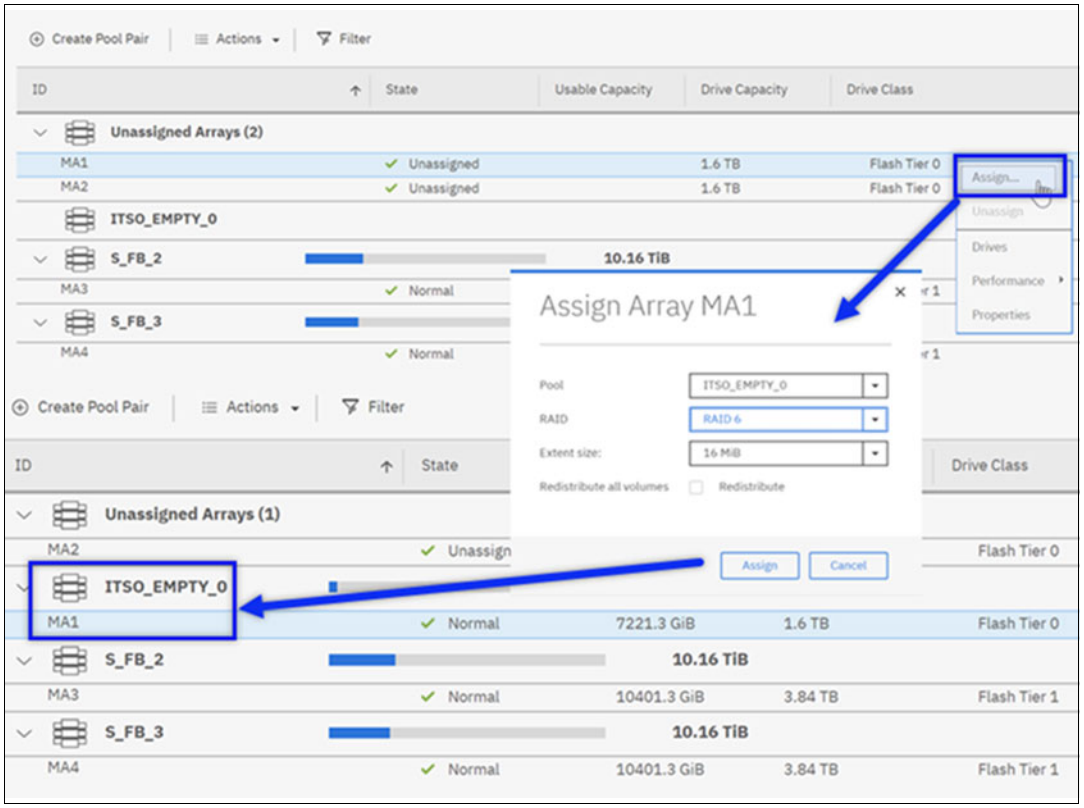


Figure 9-57 Assigning an array to a single pool

9.7.3 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 9-58 on page 275:
   - **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.

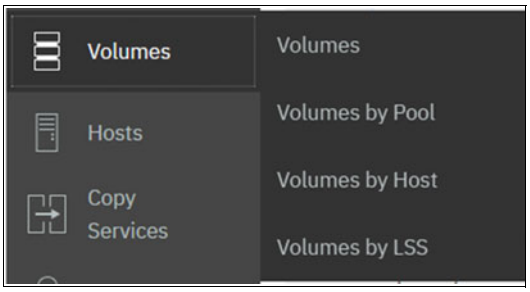


Figure 9-58 Volumes by View

2. Selecting one of the first two options opens a view that lists all the volumes or pools on the system. Figure 9-59 shows the Volumes by Pool view.

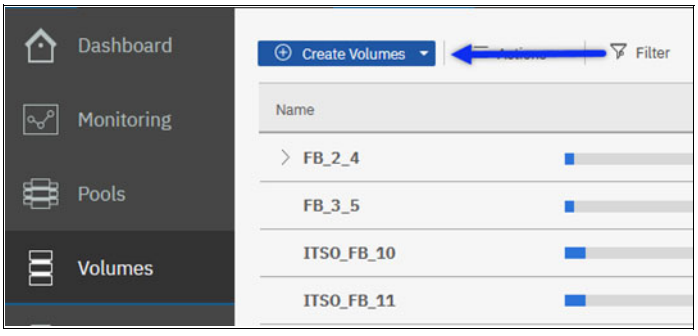


Figure 9-59 Volumes by Pool

3. From this view, click **Create Volumes**. The Create Volumes dialog opens listing various operating system platforms, as shown in Figure 9-60 on page 276.

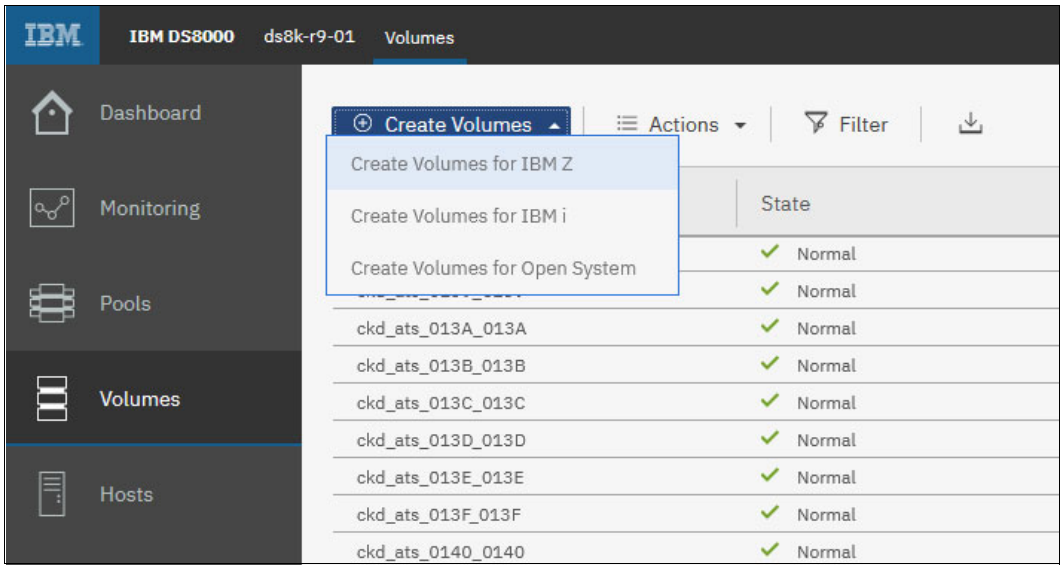


Figure 9-60 Create Volumes dialog

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 9-60 on page 276), click **Create Volumes for Open Systems**.
2. The Create Volumes for Open Systems configuration dialog for open systems hosts opens (Figure 9-61 on page 277). 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.

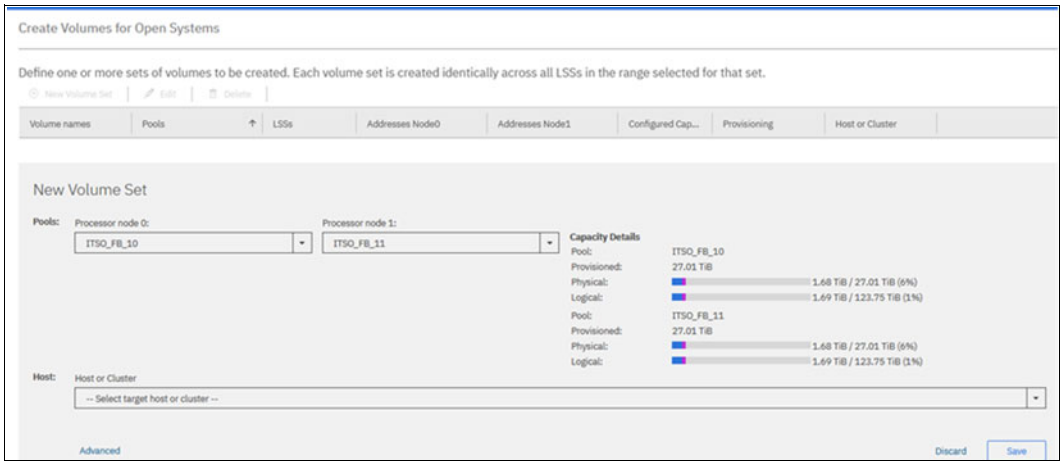


Figure 9-61 Creating open systems (FB) volumes

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.

Important: Starting with the DS8A00 G10, Standard volumes are only available when pools are made up of industry-standard NVMe drives. Those pools made up of IBM FlashCore Modules (FCM) support ESE volumes only.

The Administrator, or User with Administrator role, or storage administrator, when they create the new volumes, can assign the address range to the volume in the Advanced section, as shown in Figure 9-62 on page 278. 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, when all the volume sets are specified, click **Create** to create them in a row all at once.

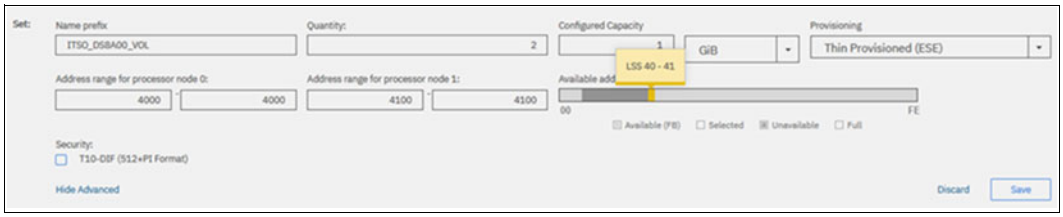


Figure 9-62 Using the Advanced section when creating volumes

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 Services 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 must both 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 9-63 on page 278 opens.

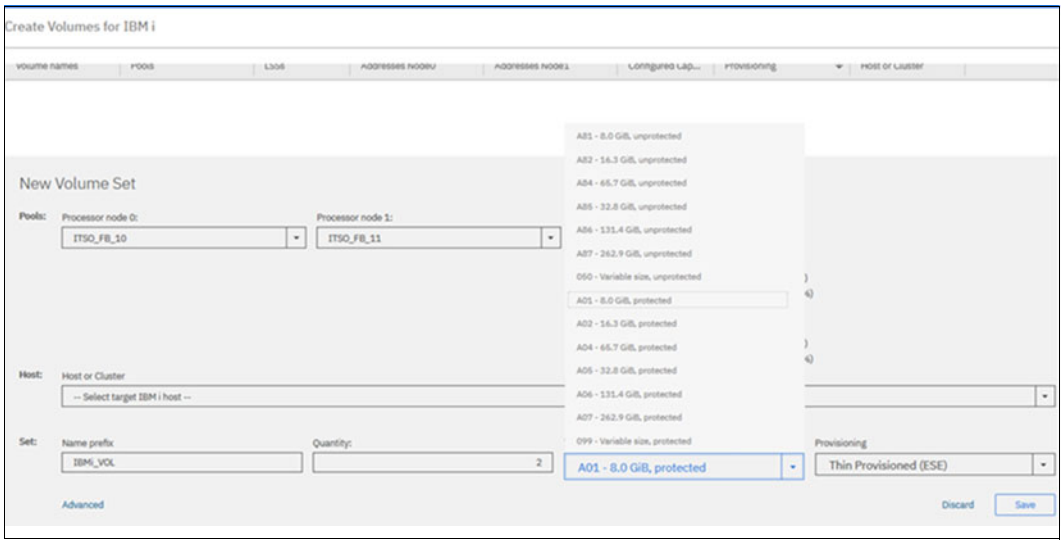


Figure 9-63 Creating IBM i volumes

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) must also 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 can specify 4-digit volume serial numbers for the volumes in the **Advanced** section. Use hexadecimal values (00–FE) for the LSSs.

Note: Dynamic Volume Expansion (DVE) of IBM i 050 and 099 volume types in increments of 1–2000 GB is available. The minimum software level of the IBM i hosts must be 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.

9.7.4 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 DS8A00, the host must be connected to the DS8A00 through the Fibre Channel Protocol (FCP).

The DS8A00 offers 4-port, 32 Gbps encryption-capable PCIe Gen3 FCP/FICON host adapters (referred to in the DS GUI as 32 Gbps). Each port can be independently configured to one of the following 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

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 9-64 on page 280).

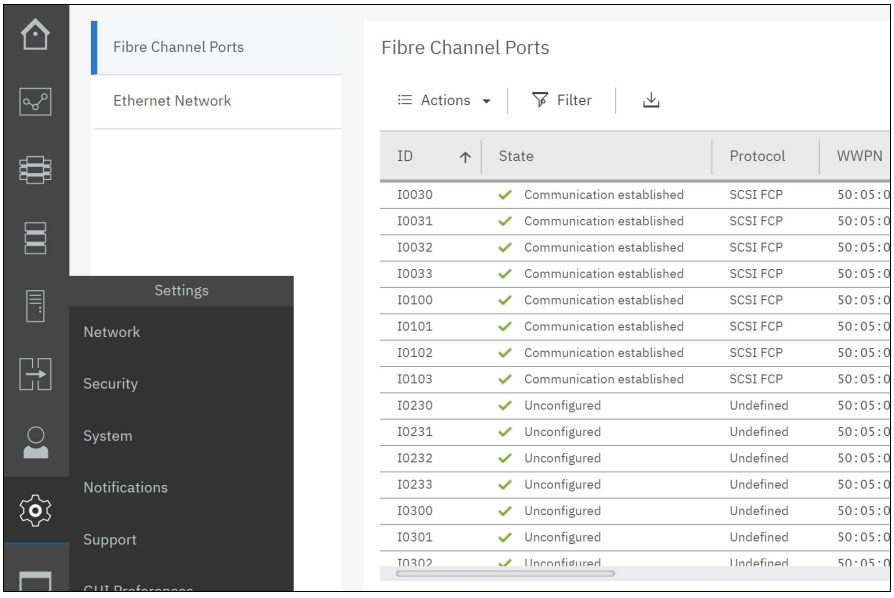


Figure 9-64 Fibre Channel Ports window

2. Select the port to modify. Multiple ports can be selected by using the Shift or Ctrl key.
3. From the Actions tab, click **Modify Protocol** to open the Modify Protocol window, as shown in Figure 9-65.

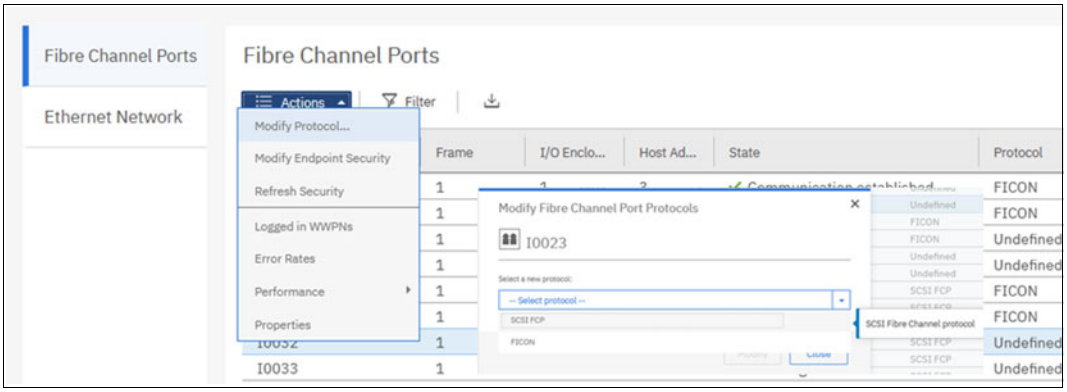


Figure 9-65 Modify FCP Protocol window

4. 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).
5. 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 DS8A00 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 DS8A00.

To configure an open system cluster, complete the following steps:

1. Create a cluster: Configure a cluster object to access the storage system.
2. Create a host: Configure a host object to access the storage system.
3. Assign hosts: Assign hosts to a cluster object.
4. 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.
5. Modify the FC port mask: Modify the FC port mask (on the DS8A00) 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 DS8A00) 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 9-66.

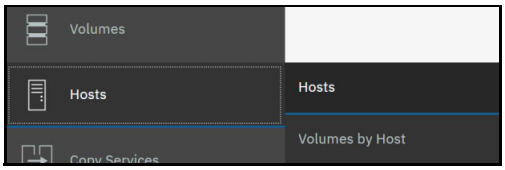


Figure 9-66 Hosts menu in the navigation

3. The Hosts window opens, as shown in Figure 9-67. Click **Create Cluster**.
4. The Create Cluster window opens, as shown in Figure 9-67. Specify the name of the cluster, and click **Create**.

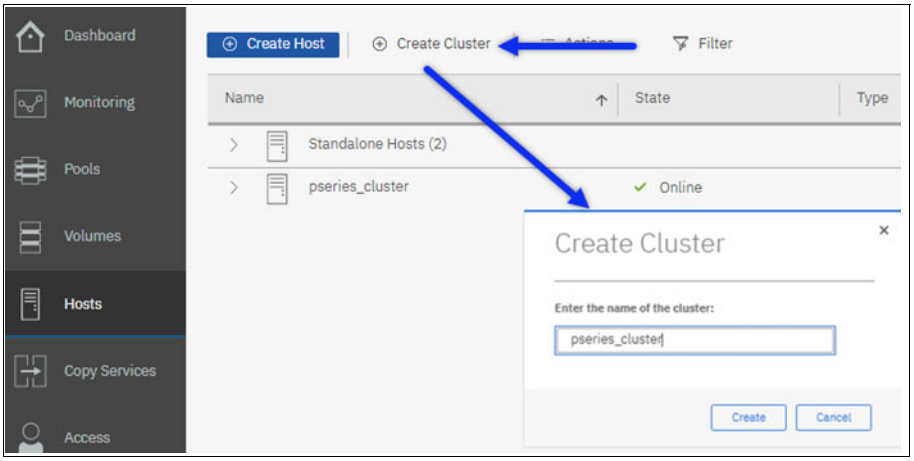


Figure 9-67 Hosts View: Create Cluster window

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 9-66 on page 281.
3. Select **Create Hosts** to assign host ports to the existing hosts or click **Create Host** to create hosts to assign the ports.
4. In the Hosts window, click **Create Host**, as shown in Figure 9-68.

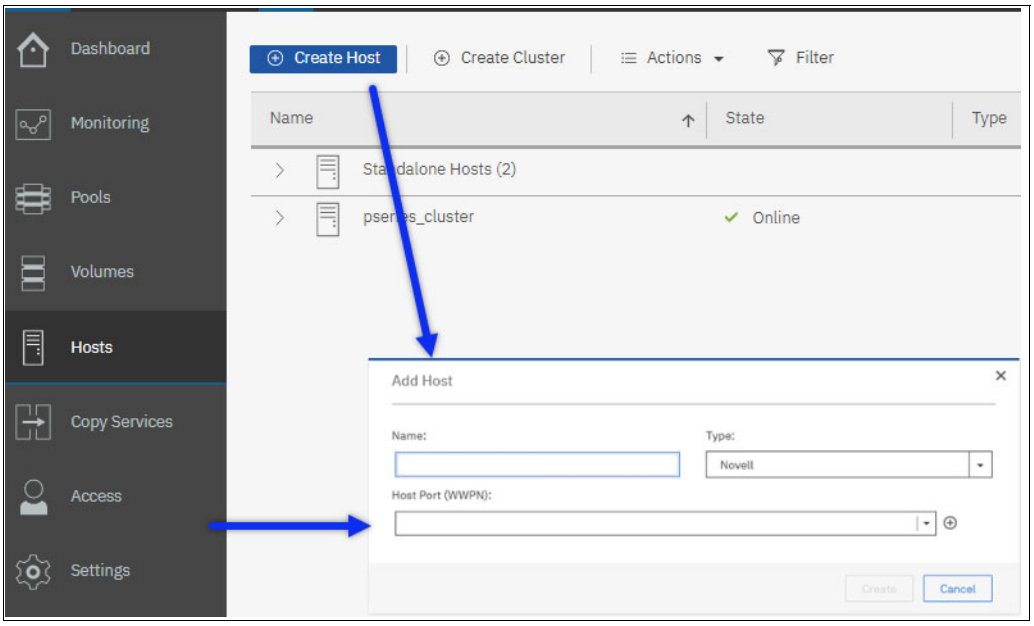


Figure 9-68 Create Host

5. The Add Hosts window opens (Figure 9-69 on page 283). 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 9-68.

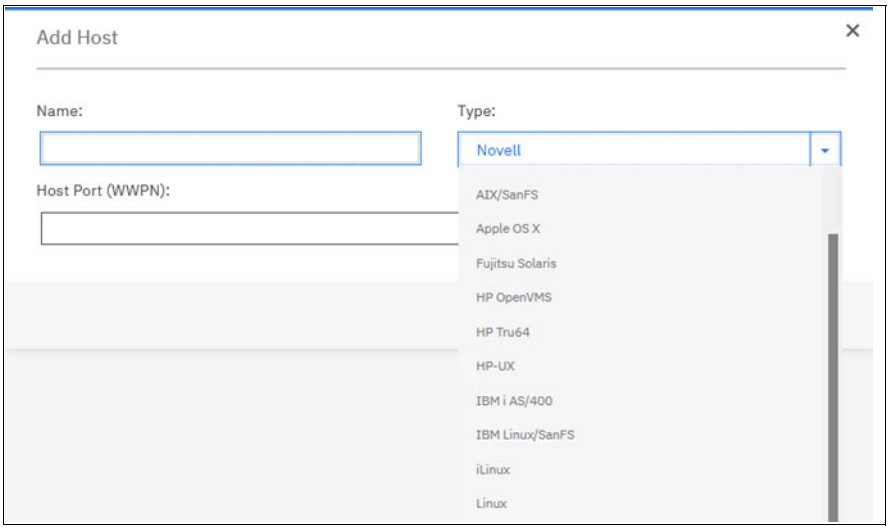


Figure 9-69 Add Host window - Host type, partial list

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 9-70.

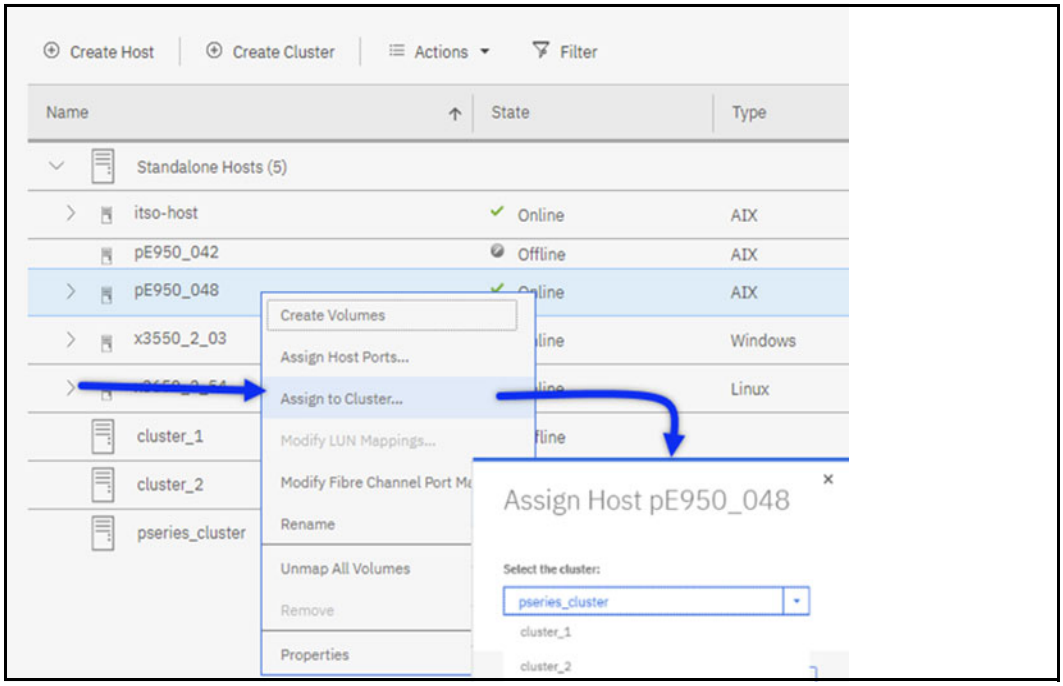


Figure 9-70 Assigning a host to a cluster

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 9-71 on page 284.

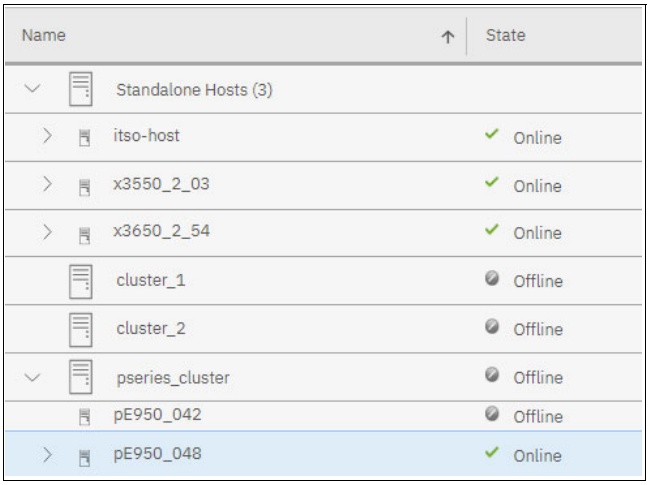


Figure 9-71 Hosts window with created cluster

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 9-72.

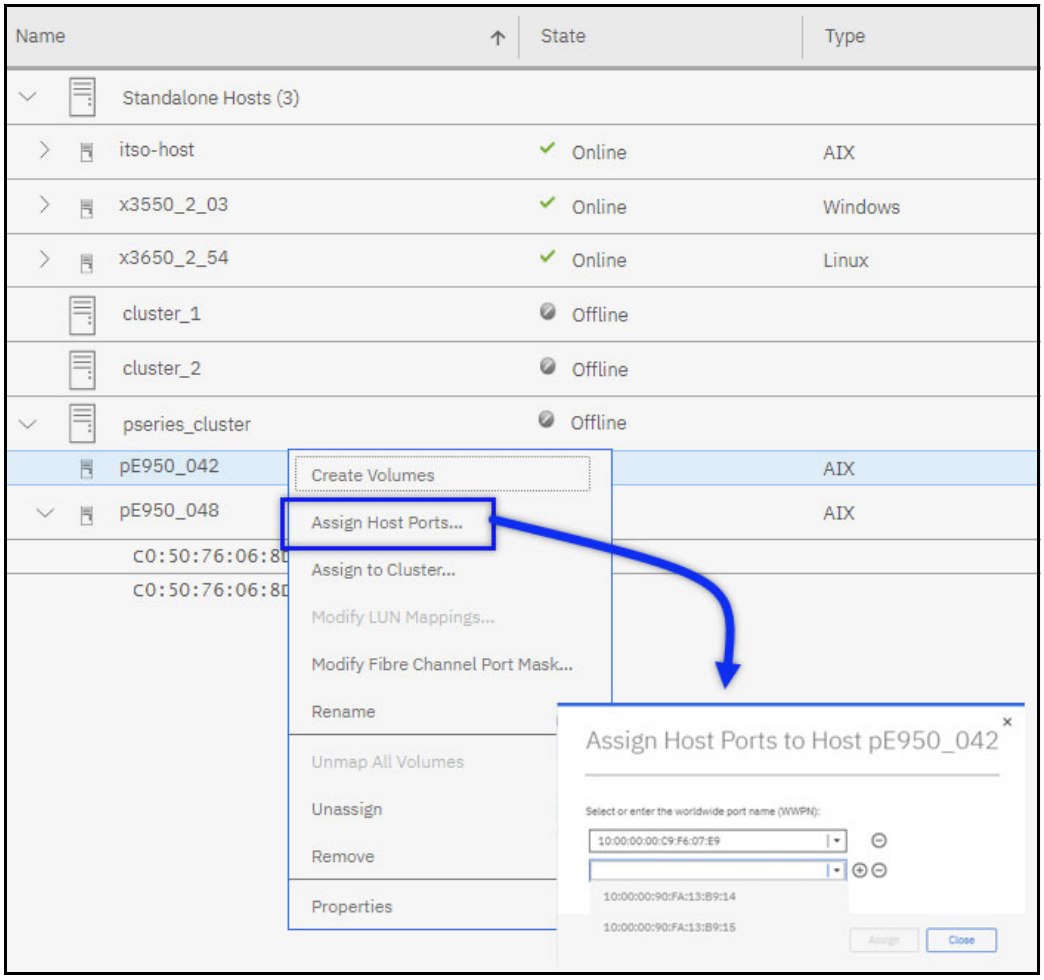


Figure 9-72 Assign Host Ports to Host

2. This action opens the Assign Host Ports to Host window.

   If the host ports are already logged in to the DS8A00, that is, the host ports are zoned so that they are visible to the DS8A00, then the available WWPNs are listed in the drop-down. If the host ports are not visible to the DS8A00, 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 9-72 on page 285. Click **Assign** to complete.

   Typically, most open system hosts have multiple FC connections to the DS8A00 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 DS8A00 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 DS8A00 system. To view them, select a host from the Hosts window, right-click, and then select **Properties**, as shown in Figure 9-73.

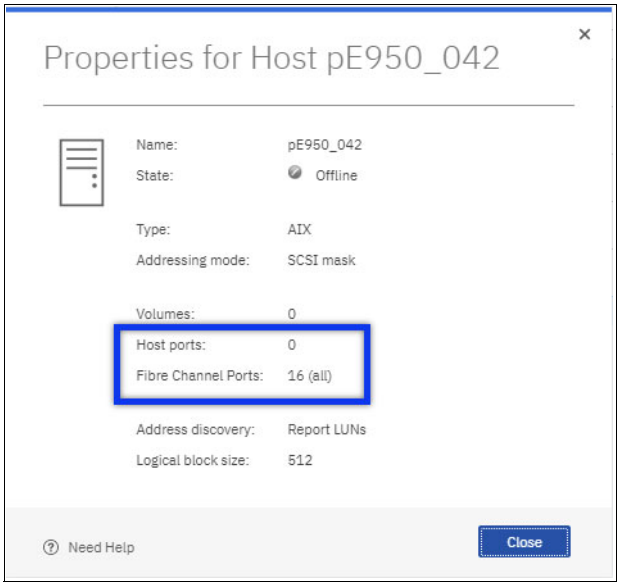


Figure 9-73 Host properties displaying the Fibre Channel port mask

If the storage 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 or from the **Actions** tab, select **Modify Fibre Channel Port Mask**. A list of the DS8A00 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 9-74 on page 287.

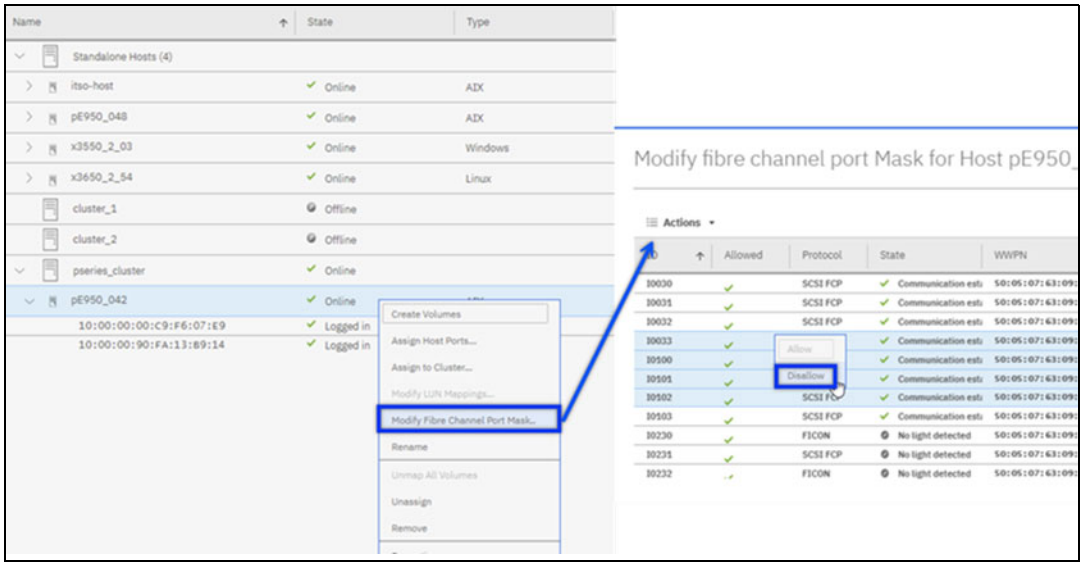


Figure 9-74 Modifying the Fibre Channel Port Mask

The properties of the selected host now reflect the number of FC ports that have access as shown in Figure 9-75.

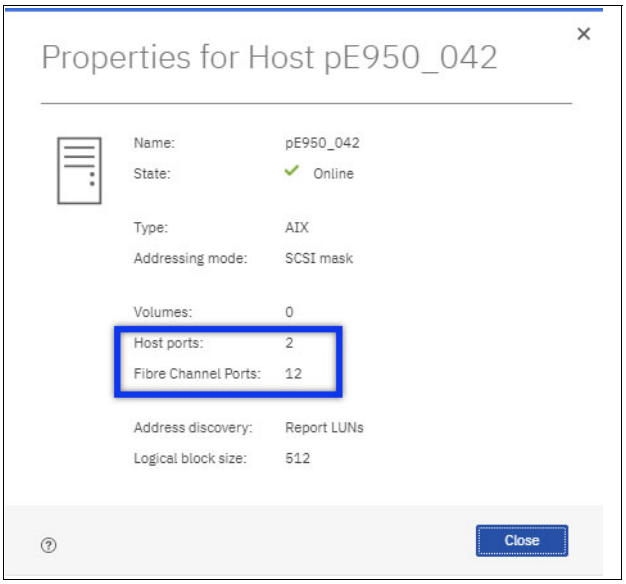


Figure 9-75 Host properties with FC port masking

9.7.5 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 9-76 on page 288. From the Actions tab, or by right-clicking, select **Map to Host or Cluster**.

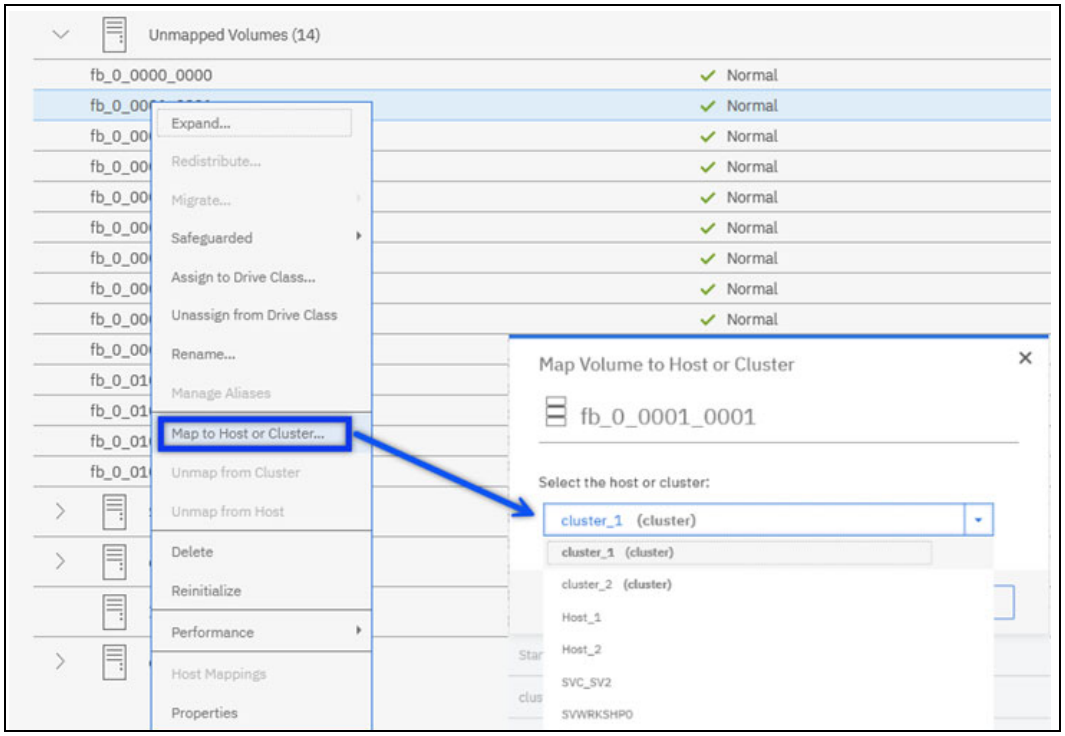
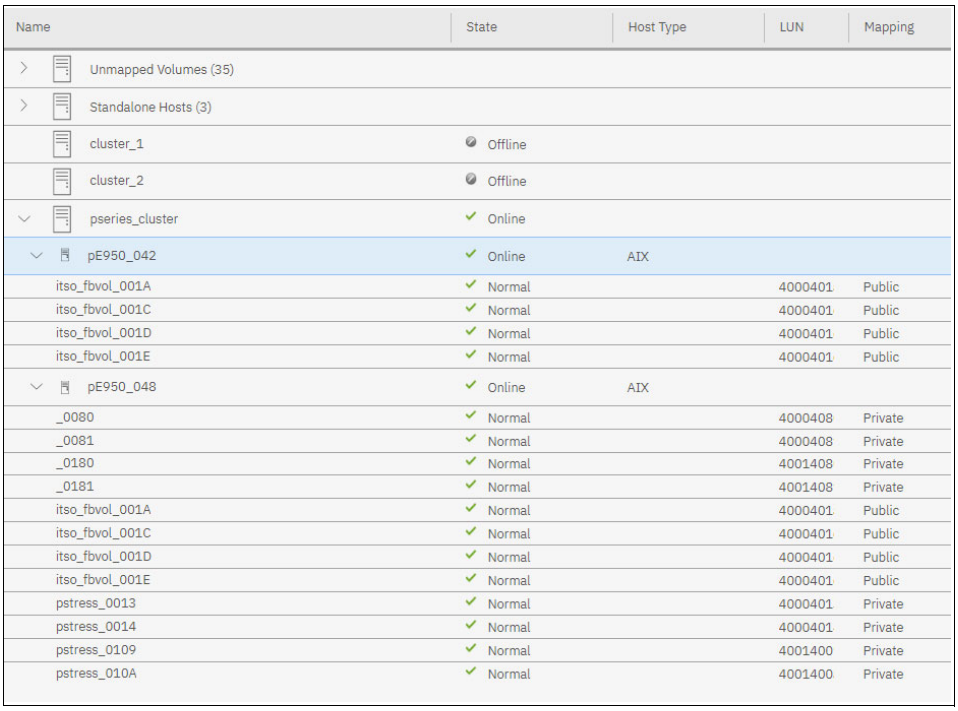


Figure 9-76 Mapping unmapped volumes to a host or cluster

3. 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 9-76, 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 9-77 on page 289 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 9-77 Mapping volumes to a cluster

9.8 Logical configuration for Count Key Data volumes

This section describes the logical configuration for CKD volumes for IBM Z.

9.8.1 Configuration flow

As with the FB configuration, logical configuration for CKD volumes is simplified and can now be accomplished with a few steps:

1. Create a CKD pool pair and assign arrays to the pools.
2. Create the CKD LSSs (also known as Logical Control Unit (LCU)).
3. Create the CKD volumes.
4. Configure the alias volumes (Parallel Access Volumes (PAVs)).
5. Configure the FC ports for FICON.

9.8.2 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. The Create Pool Pair window is opened as seen in Figure 9-78 on page 290.

4. Select the **Storage Type.** In this case Count Key Data (CKD) is preferred.
5. Decide whether the pair will be Compressed. Compression is important:
   a. If no compression is selected, then industry-standard NVMe disks are used. Standard, full-provisioned, volumes can be created. No arrays are automatically assigned. See Figure 9-49 on page 271 and Figure 9-50 on page 271 as an example.
   b. If compression is selected, FCMs are used. Only ESE volumes can be created. Arrays are automatically assigned during pool pair creation
   c. Assign a name to the pool pair. This name is used as the prefix for the pool pair ID.
6. Choose the Extent size. Small extents of 21 cylinders, or large extents of 1113 cylinders.
7. Select the number of arrays to assign to the pool pair.

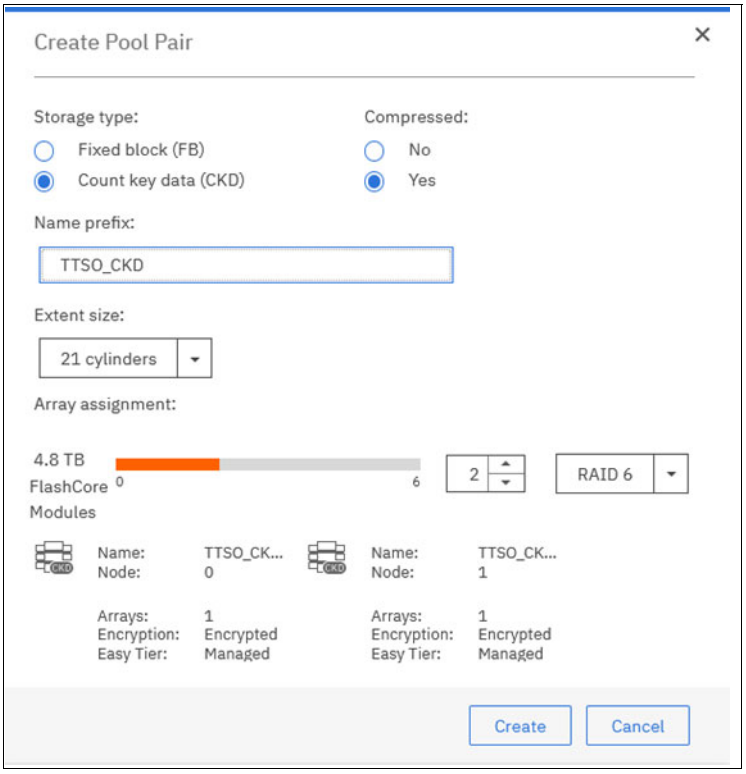


Figure 9-78 CKD pool pair creation and assigning arrays to the pool pair

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.

9. When the pool pair creation is complete, the arrays are assigned to the pool pair as shown in Figure 9-79 on page 291. The DS GUI configures the selected arrays for CKD storage and distributes them evenly between the two pools.

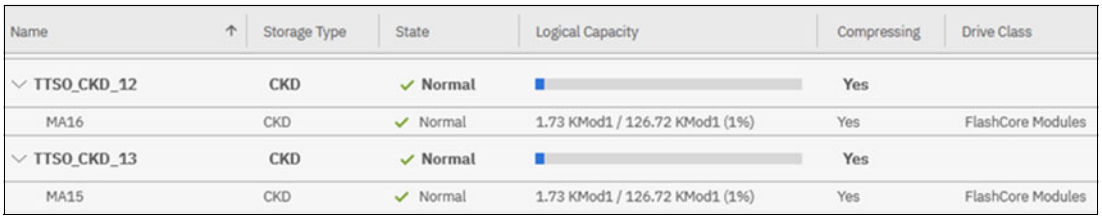


Figure 9-79 CKD pool pair after creation and assignment of arrays

Note: With the added resiliency of RAID 6 and RAID 10, RAID 5 is no longer supported on DS8A00 models.

For more information about the supported drive types and available RAID levels for DS8A00 models, see Chapter 4, “IBM DS8A00 reliability, availability, and serviceability” on page 109.

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 9-80 shows the Arrays by Pool window, which shows how to assign the arrays.

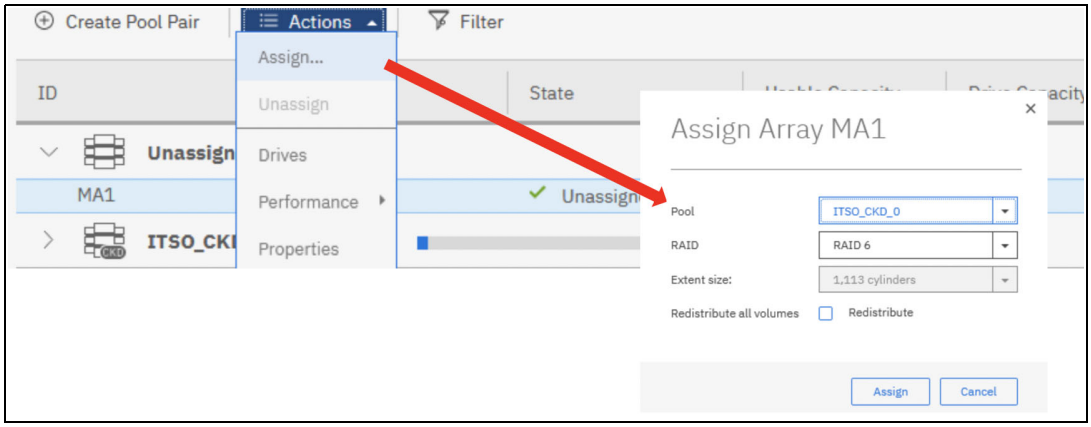


Figure 9-80 Manually assigning arrays to an existing pool

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 273.

9.8.3 Creating CKD logical subsystems

Note: You can create LSS ranges, exact volume address ranges, and aliases in one step. For an example, see 9.8.4, “Creating CKD volumes” on page 295.

The DS8000 LSS emulates a CKD Logical 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 9-81.

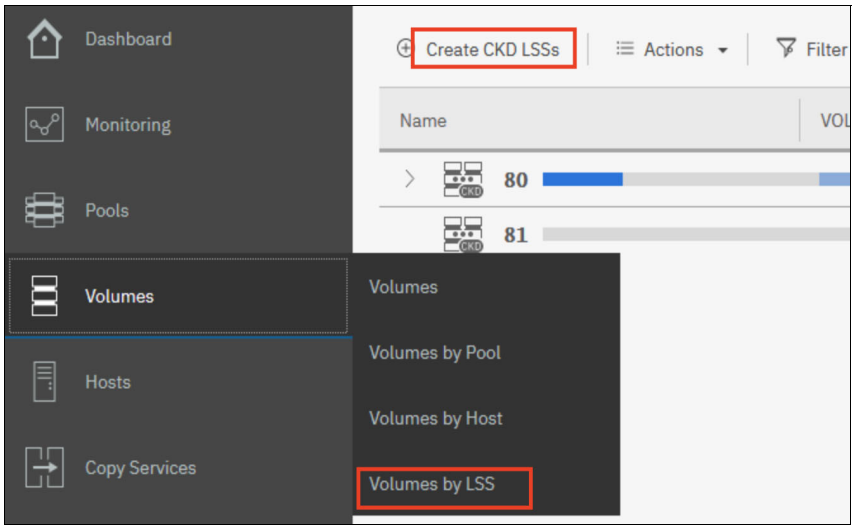


Figure 9-81 Creating CKD volumes by LSS

4. The Create CKD LSSs window opens, as shown in Figure 9-82 on page 293. 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.

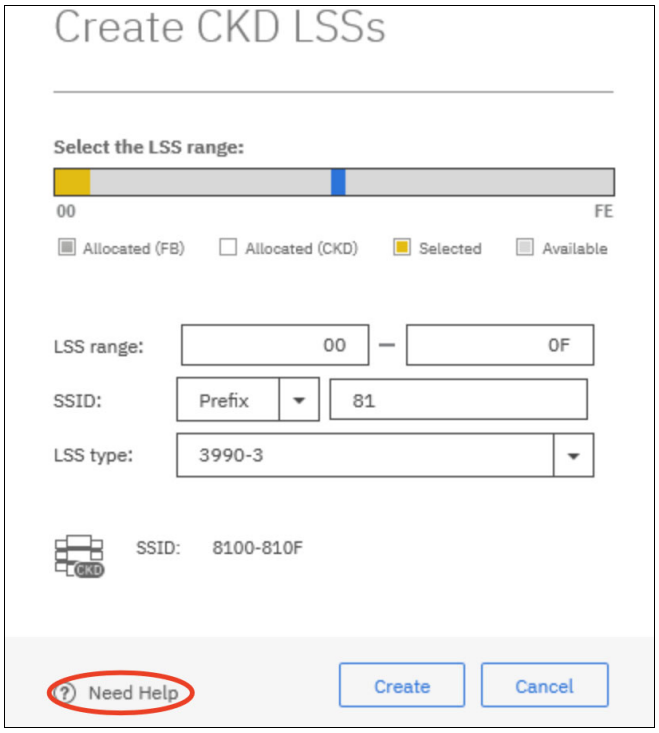


Figure 9-82 Defining CKD LSSs

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 9-83 shows the 16 LSSs that are created.



Figure 9-83 CKD LSSs that are created

5. 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 9-84.

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 as they must be unique in an IBM Z environment.

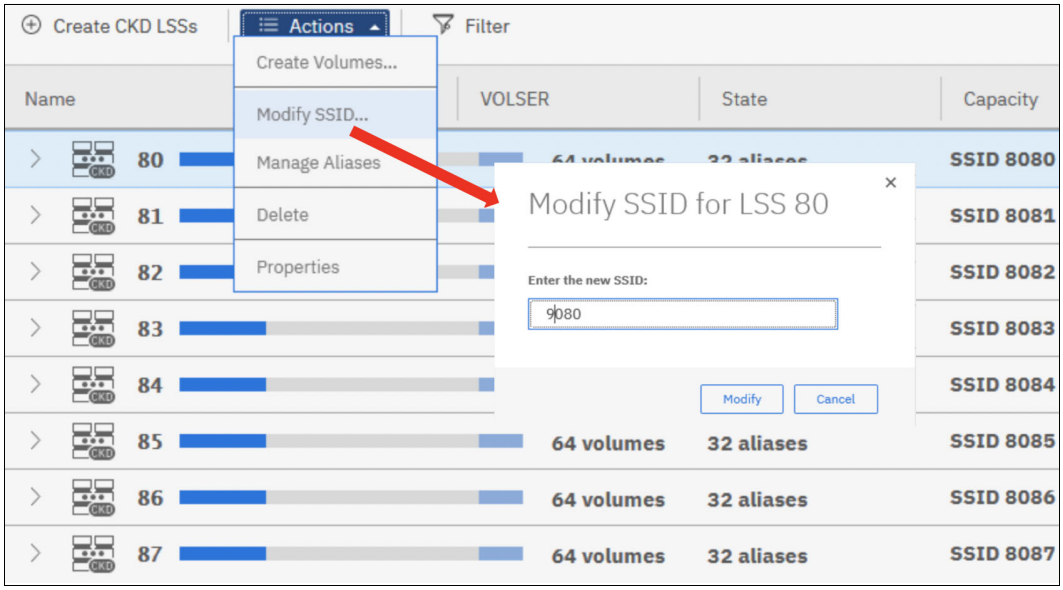


Figure 9-84 Modifying the CKD LSS SSID

Important: Use caution when changing SSIDs that are in a Copy Services relationship. If a change must be done, then the entire relationship must be broken down including the PPRC paths.

SSIDs must be unique across all DS8000 systems participating in copy services relationships.

Notes:

- Occasionally, the DS8A00 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 267.
- SSIDs must be unique across all DS8000 systems participating in copy services relationships.

9.8.4 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 9-85 on page 295.

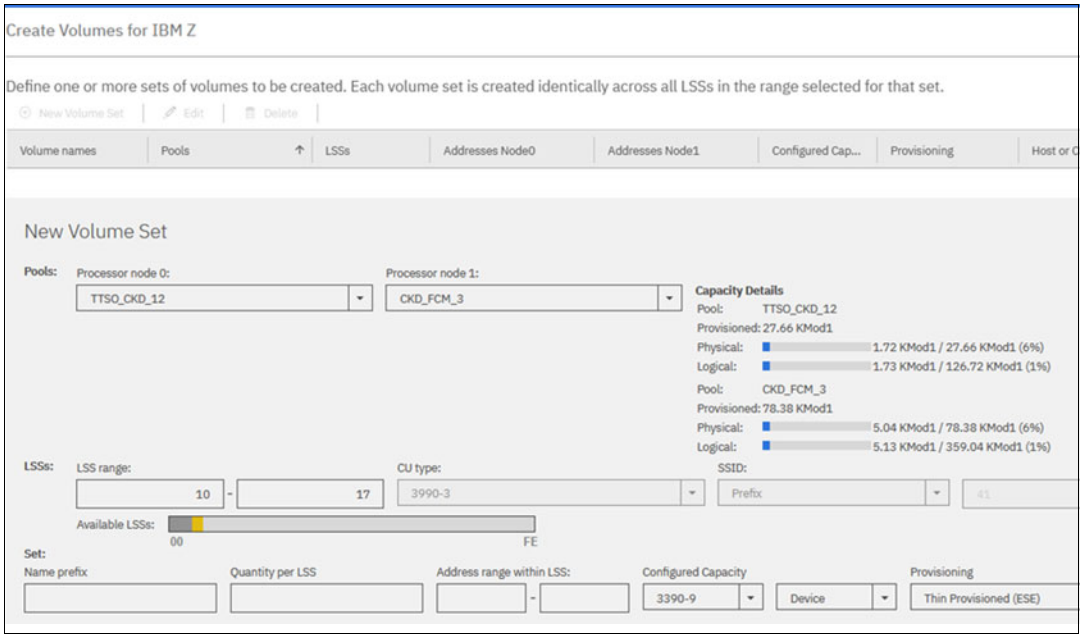


Figure 9-85 Creating multiple groups of CKD volumes

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)**.

> **Note:** Pools with FCM storage support only Thin Provisioned (ESE) volumes. To configure “Standard” volumes, select or create a pool with Industry-standard NVMe drives.

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.

9.8.5 Creating CKD parallel access volumes

The DS8A00 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). Following that is now HyperPAV which simplifies operational procedures by removing static alias to base bindings and associating a PAV alias device to a pool. There is also SuperPAV mode which supports the borrowing of aliases from neighboring LSSs, thus requiring fewer aliases to be defined per system.

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 are typically configured to start at FF and decrement (decrease) toward 00. A storage 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 per 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 DS8000 and IBM Z Synergy*, 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 9-86 on page 297.

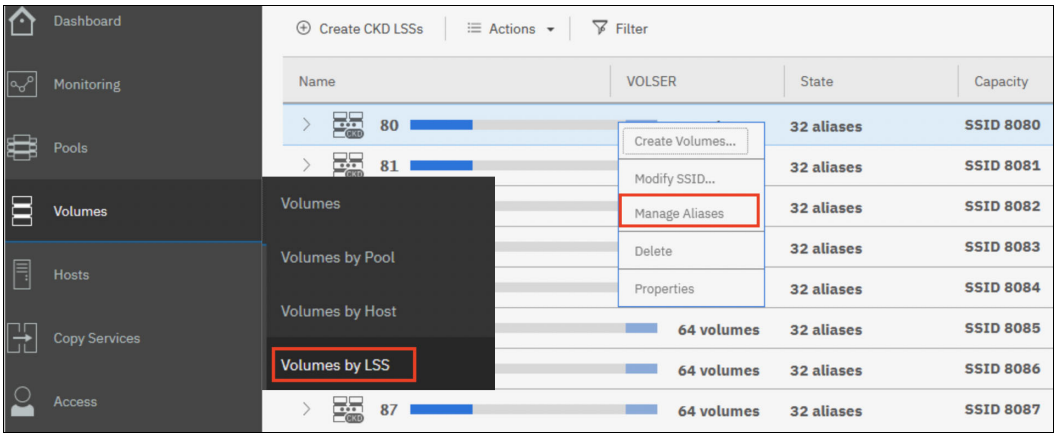


Figure 9-86 Creating aliases for the LSS

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 9-87. Enter the number of aliases to create. The example in Figure 9-87 shows 32 aliases being created for LSS 80.

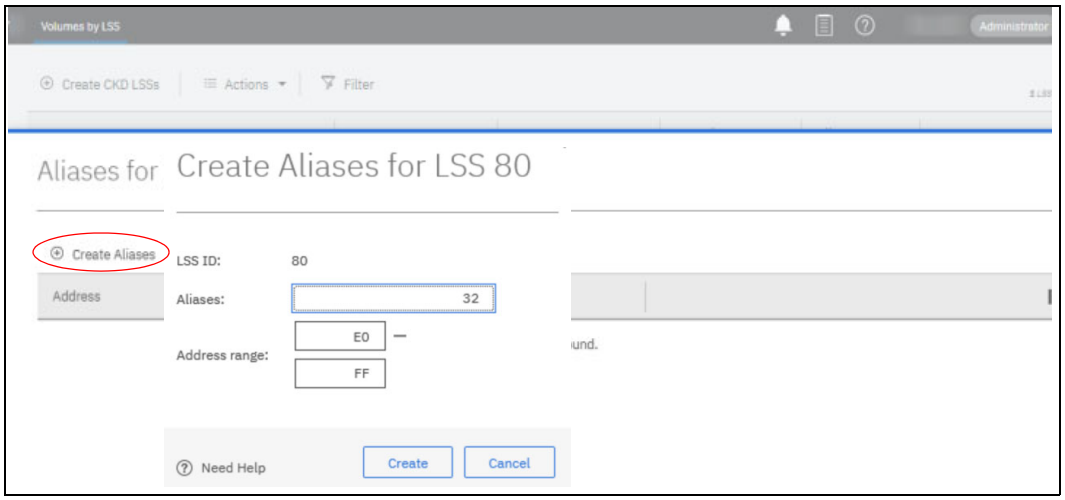


Figure 9-87 Creating 32 aliases for LSS 80

4. Click **Create**. The aliases are created for LSS 80, as shown in Figure 9-88.



Figure 9-88 Thirty-two aliases are created for LSS 80

5. 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 9-89 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 9-89, the 32 aliases are created against the lowest base volume address 8000 (`ITSO_CKD_8000`).

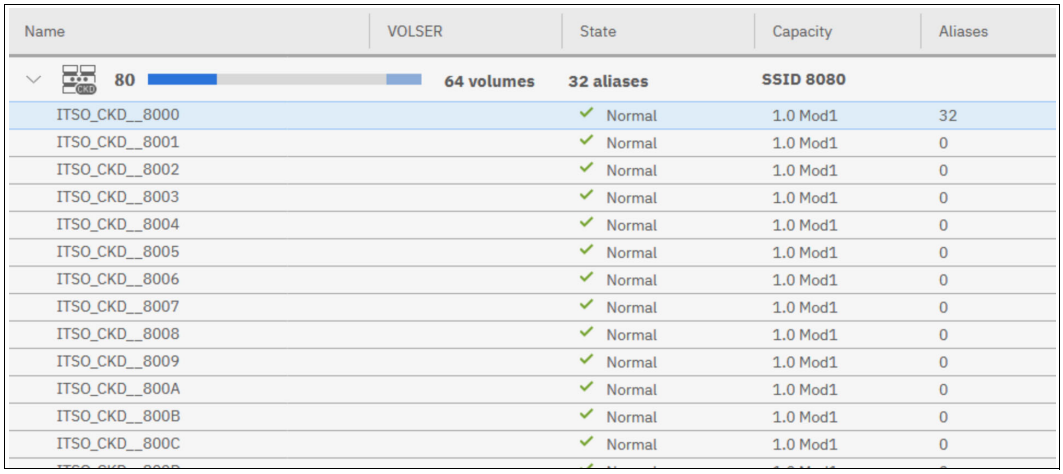


Figure 9-89 Thirty-two aliases against the lowest base volume address

6. 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 shown in Figure 9-90.

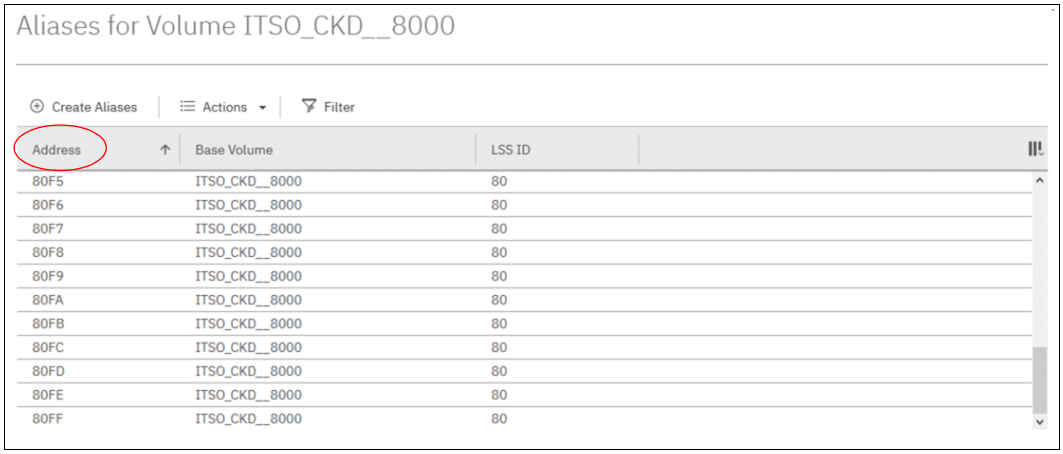


Figure 9-90 List of aliases with their alias IDs starting at F5 and in ascending order

Note: The alias IDs start at F5 and are in ascending order, as shown in Figure 9-90.

7. 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 9-91.

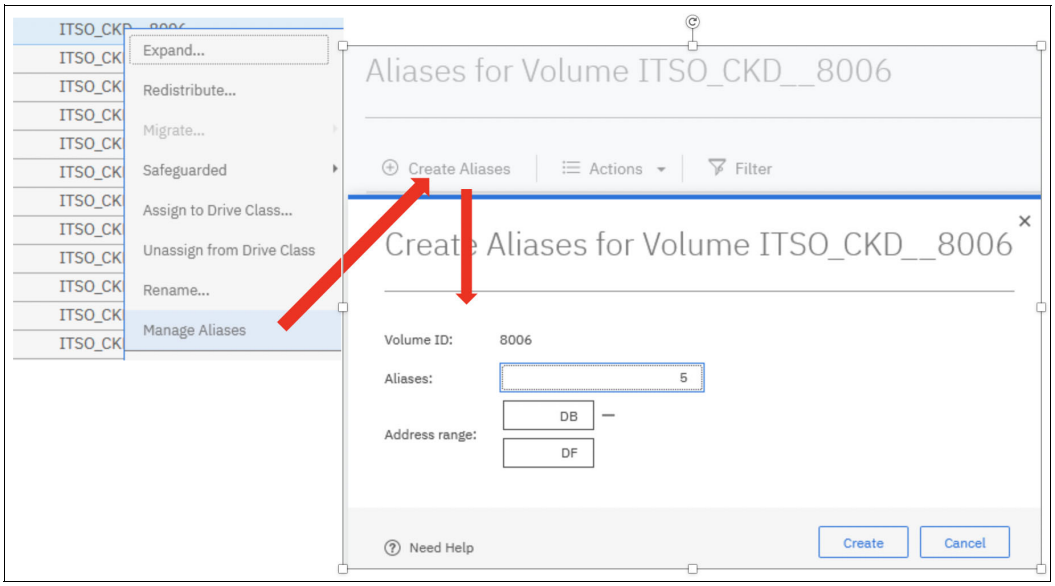


Figure 9-91 Configuring aliases for a single base volume

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 ascending order, as shown in Figure 9-92. (Aliases E0–FF were created before).



Figure 9-92 List of five aliases that are created for a single base address

9.8.6 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 DS8A00 FICON.

To set the FC port protocols of the FC ports that the host uses to communicate with the DS8A00, complete these steps:

1. Select **Settings** → **Network** → **Fibre Channel Ports**. Select a port 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 9-93.

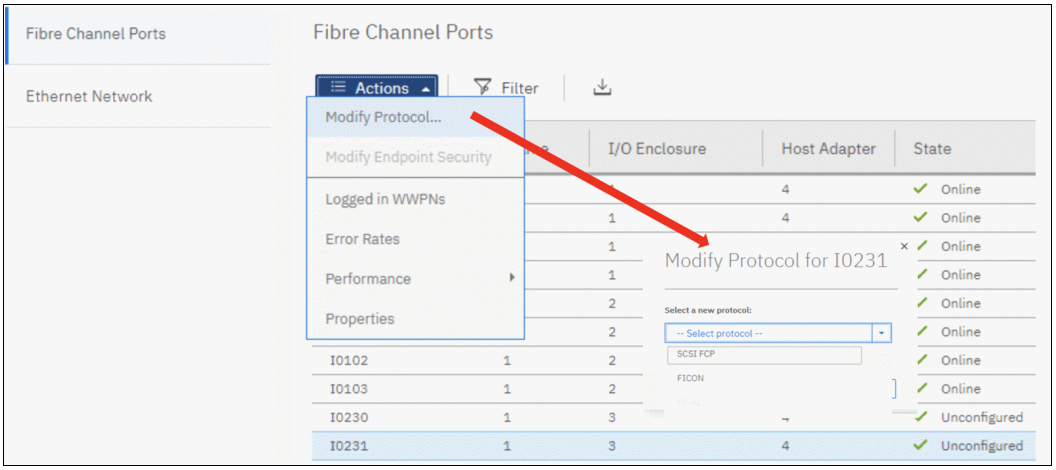


Figure 9-93 Modifying the FC port protocols

9.9 Expanding volumes

The DS8A00 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 (Figure 9-94 on page 301):

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 might appear to inform 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.

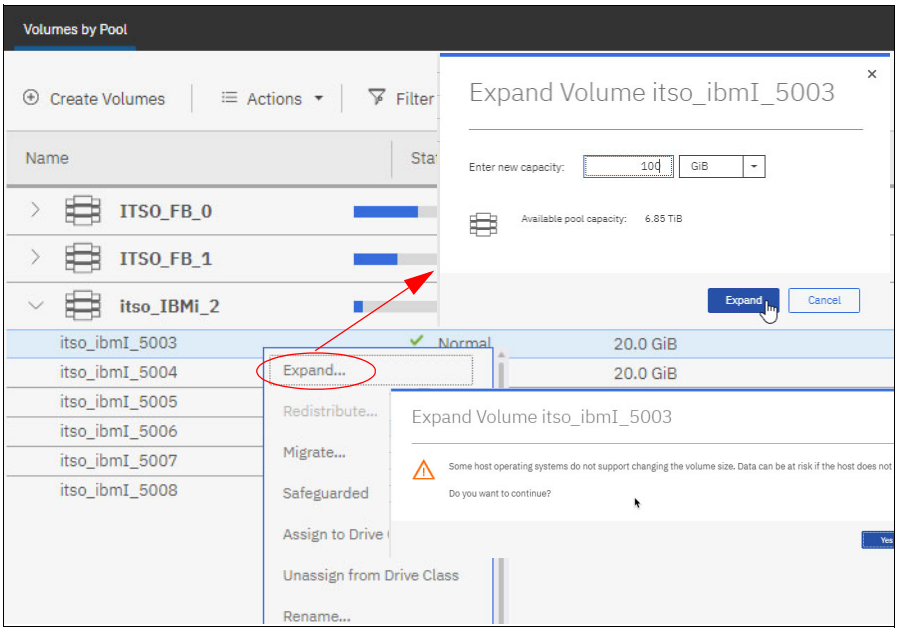


Figure 9-94 Expanding a volume

The storage administrator can also **Expand** the Safeguarded capacity, as shown in Figure 9-95. For more information, see *IBM Storage DS8000 Safeguarded Copy*, REDP-5506.

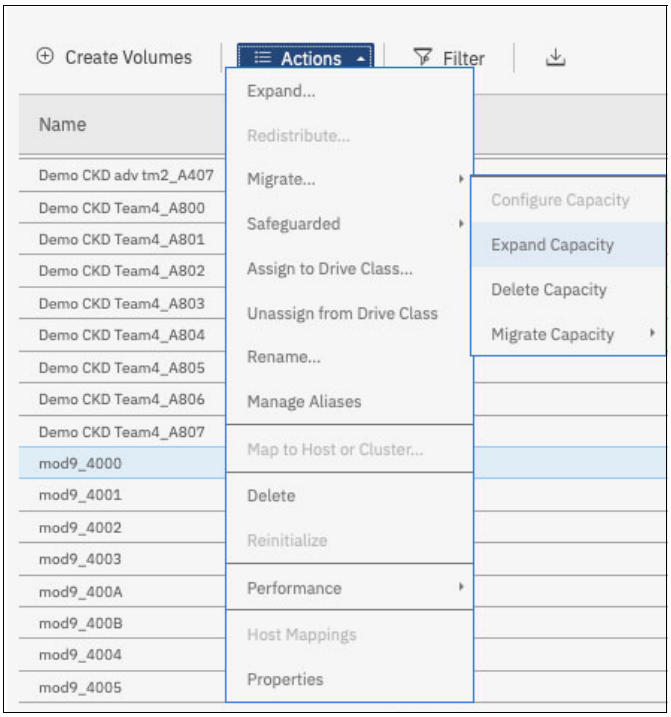


Figure 9-95 Expanding Safeguarded capacity

9.10 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 9-96.

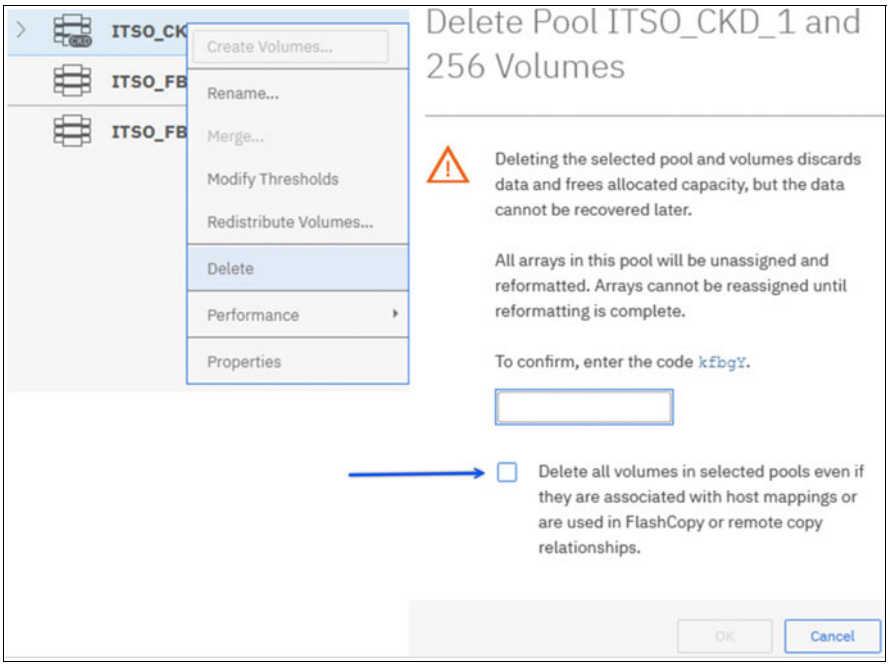


Figure 9-96 Deleting a pool with volumes

9.11 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*.

For FB volumes, a volume is considered to be in use if it is in a Copy Services 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 Copy Services 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 9-97 on page 303.

Important: 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 9-97 on page 303.

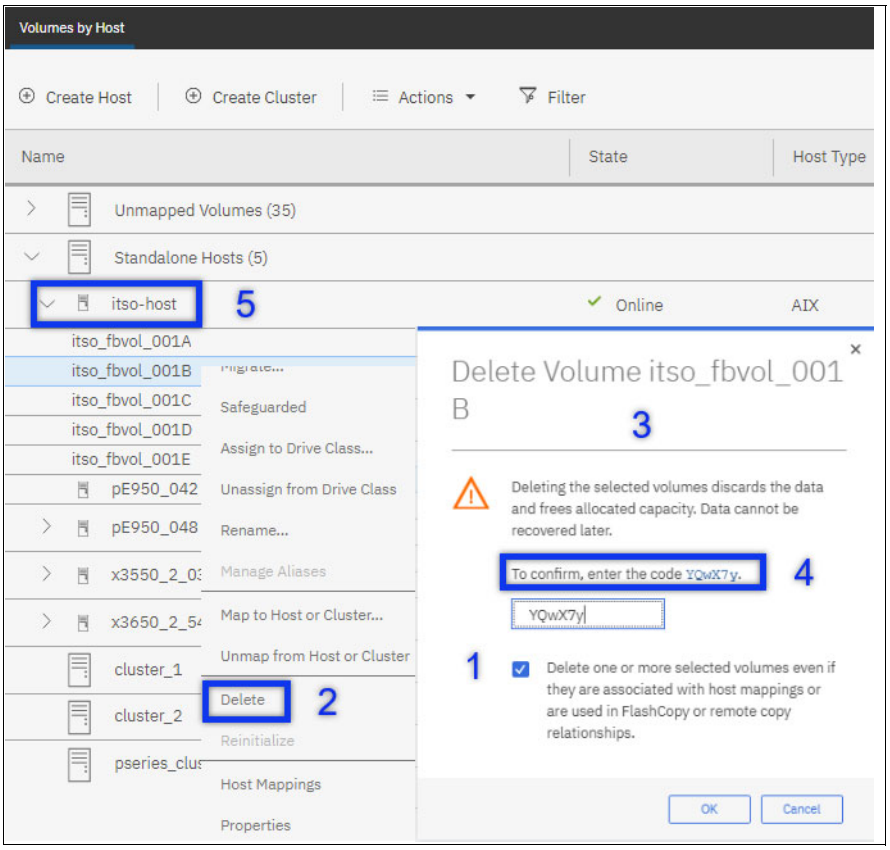


Figure 9-97 Deleting volumes

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 is completed.

9.12 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 9-98.

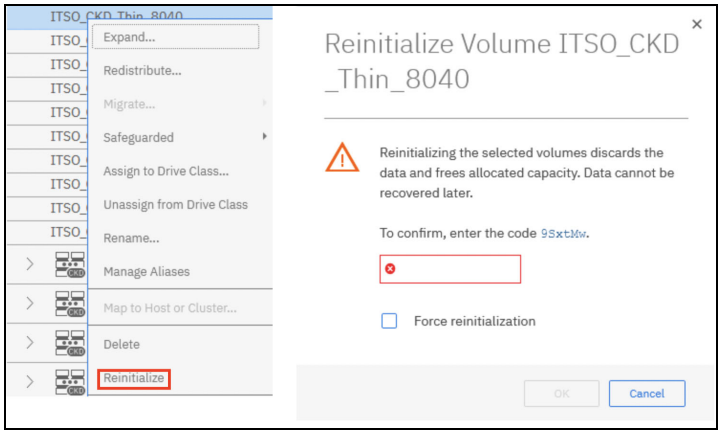


Figure 9-98 Reinitializing an online ESE volume

9.13 Easy Tier support

You can assign the volume to an Easy Tier drive class, or you can specify a tier that the volume cannot use.

Assigning a volume to a drive class

Select a volume, and then click **Assign to Drive Class**, as shown in Figure 9-99 on page 305. A window with options that are based on the drive classes that are in the pool opens. The following options are available, depending on the drive classes that are in the pool:

- Performance Flash Tier 1 (Industry-standard NVMe) - 1.6 TB
- Capacity Flash Tier 2 (industry-standard NVMe) - 1.92 TB, 3.84 TB, and 7.68 TB
- Compression Flash Tier (IBM FCM4) - 4.8 TB, 9.6 TB, and 19.2 TB

There are three drive classes available as noted previously. You cannot change the drive class of a device. Figure 9-99 on page 305 shows examples of two. “**A**” denotes assignment of Performance Flash Tier 1, and “**B**” shows assignment of Compression Flash Tier.

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 258.

For more information about Easy Tier, see *IBM DS8000 Easy Tier*, REDP-4667.

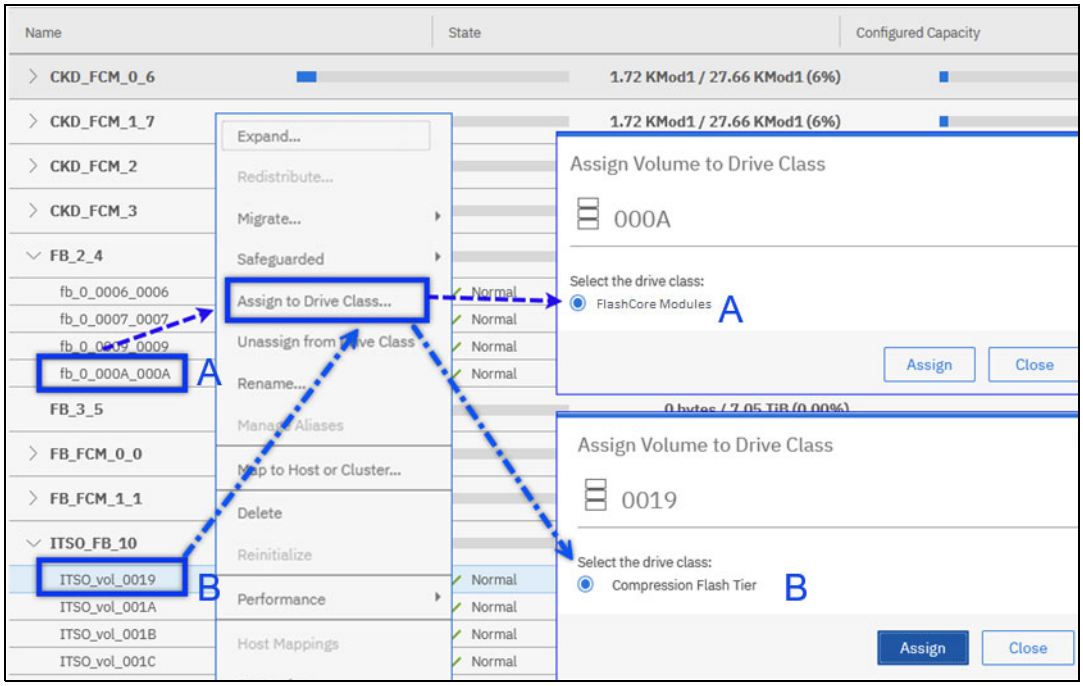


Figure 9-99 Assigning a volume to a drive class

9.14 Monitoring system health

The DS8A00 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 DS8A00 storage system. It displays system hardware states, and also provides access to actions that can be performed on the hardware components.

Figure 9-100 shows an example of the overall status of a system as displayed by the home dashboard.

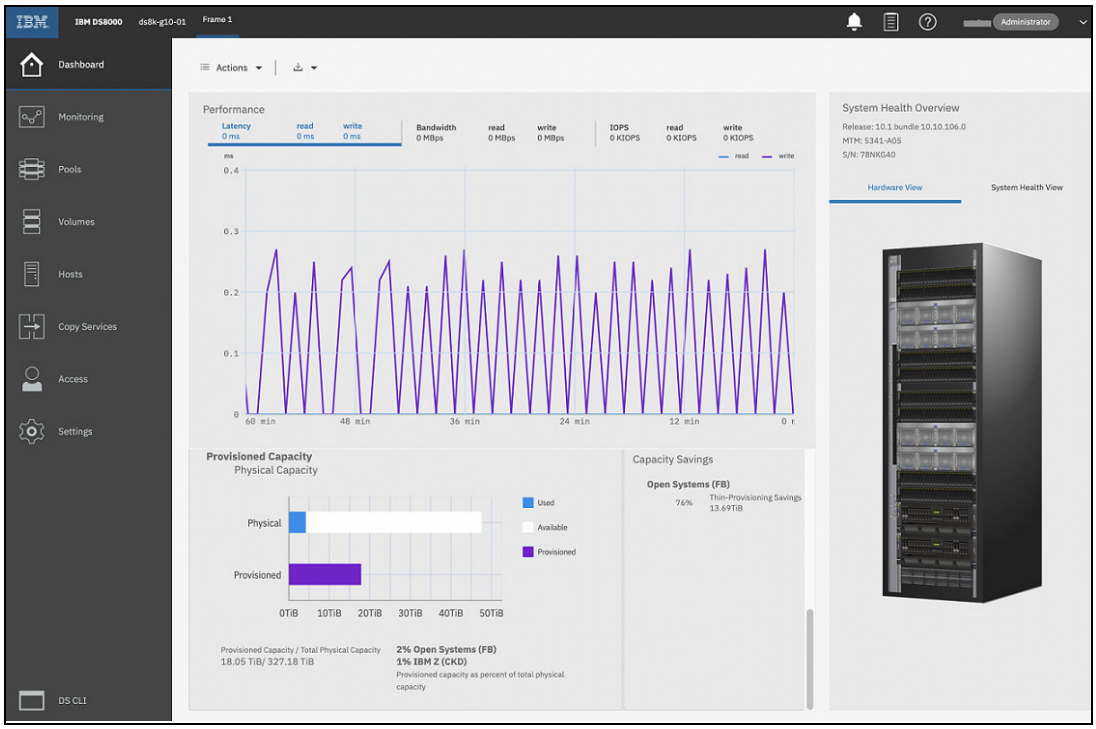


Figure 9-100 System Monitoring Dashboard

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.

- Capacity Savings

  Displays the savings from Thin-Provisioning, Compression, and Capacity for both Open Systems and IBM Z (CKD) environment. This is shown for the Compression, and the Non-Compression pools

- 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 DS8A00 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 9-101). Click the **Alerts** icon to see the alerts. Click the specific alert to view the corresponding event in the Events window.



Figure 9-101 Dashboard Events

9.14.1 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.

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 can also display the system health overview by clicking the **System Health View** icon, as shown in Figure 9-102.

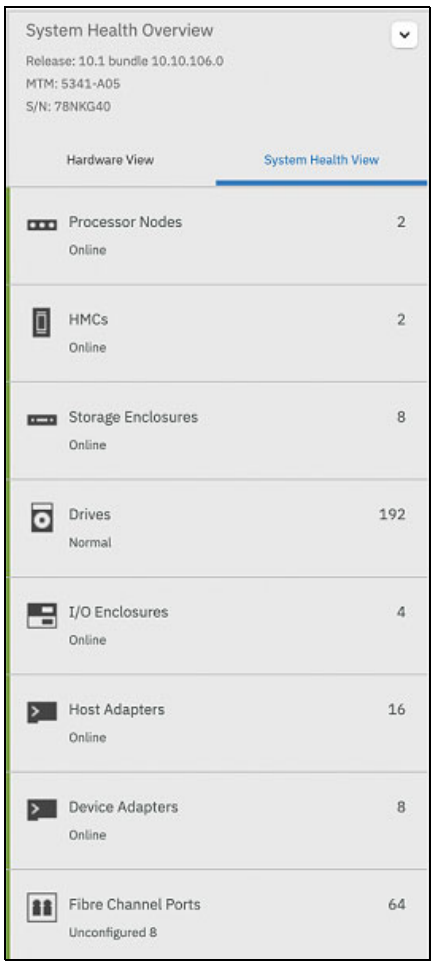


Figure 9-102 System Health Overview

For more information about the processor nodes, click **Processor Nodes**.

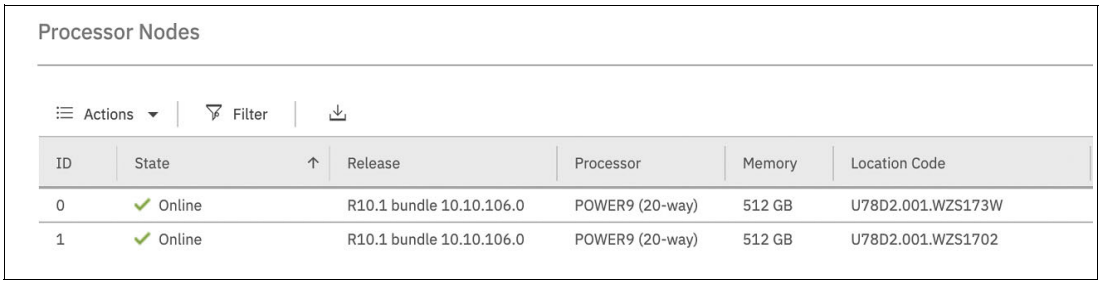


Figure 9-103 Processor Nodes - detailed view

Here are the node attributes that are shown in Figure 9-103:

- 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 installed 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 DS8A00 has two HMCs. The HMC provides a standard interface on a dedicated console to control managed systems. Point to the HMC component (Figure 9-102 on page 308) and click **HMCs** to display the detailed attributes for both HMCs, as shown in Figure 9-104.

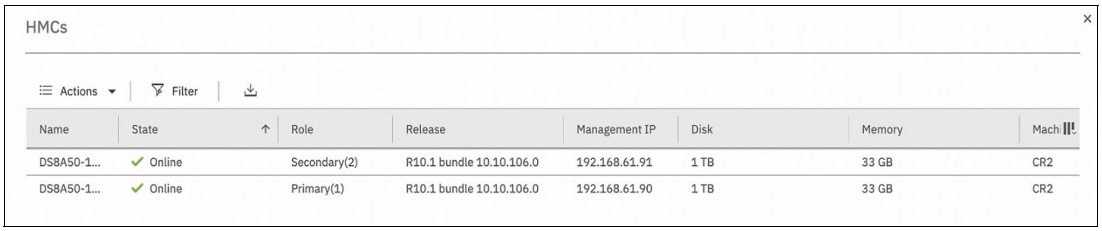


Figure 9-104 HMCs - detailed view

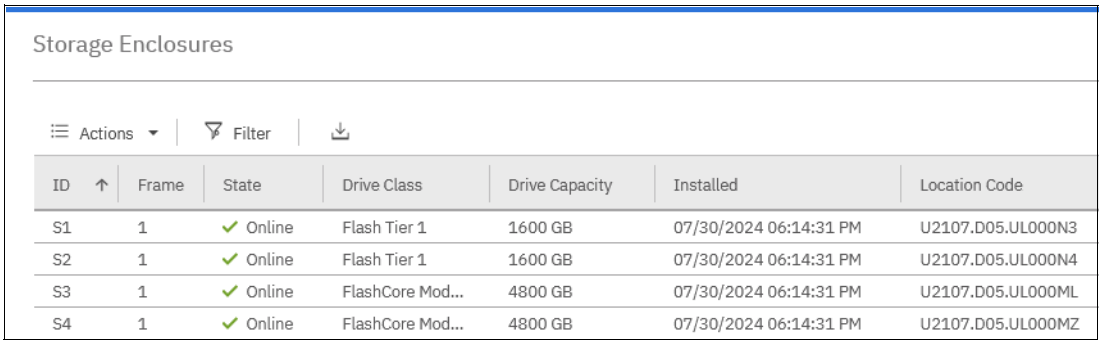
These are some of the attributes that are displayed for the HMC component in Figure 9-104 on page 309:

- 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 DS8A00 storage system. The storage enclosure also provides a mechanism to allow the drives to communicate with one or more host systems. All enclosures in the DS8A00 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 9-102 on page 308. Some of the attributes of the storage enclosure are shown in Figure 9-105.



ID	Frame	State	Drive Class	Drive Capacity	Installed	Location Code
S1	1	Online	Flash Tier 1	1600 GB	07/30/2024 06:14:31 PM	U2107.D05.UL000N3
S2	1	Online	Flash Tier 1	1600 GB	07/30/2024 06:14:31 PM	U2107.D05.UL000N4
S3	1	Online	FlashCore Mod...	4800 GB	07/30/2024 06:14:31 PM	U2107.D05.UL000ML
S4	1	Online	FlashCore Mod...	4800 GB	07/30/2024 06:14:31 PM	U2107.D05.UL000MZ

Figure 9-105 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 Performance Flash Tier 1, Capacity Flash Tier 2, and Compression FlashCore Module 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 Performance Flash Tier 1, Capacity Flash Tier 2, or Compression FlashCore Module 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 9-100 on page 306). This action shows the storage enclosure and installed storage devices in more detail (Figure 9-106). You can also select **Drives** from the System Health Overview to display information for all installed drives.

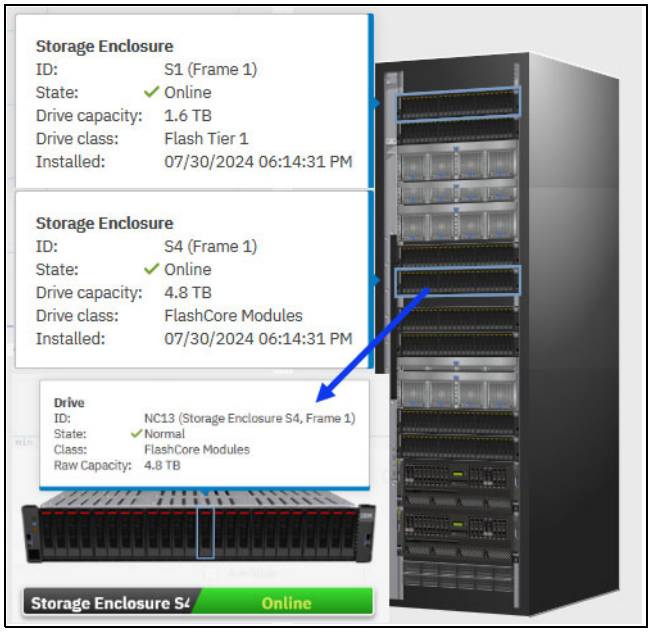


Figure 9-106 Drive detail view in a storage enclosure

The attributes for the drives in Figure 9-106 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.
- Class: The drive class.
- Capacity: The raw capacity of the drive.

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 9-100 on page 306) 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 9-102 on page 308) to open the dialog that is shown in Figure 9-107.

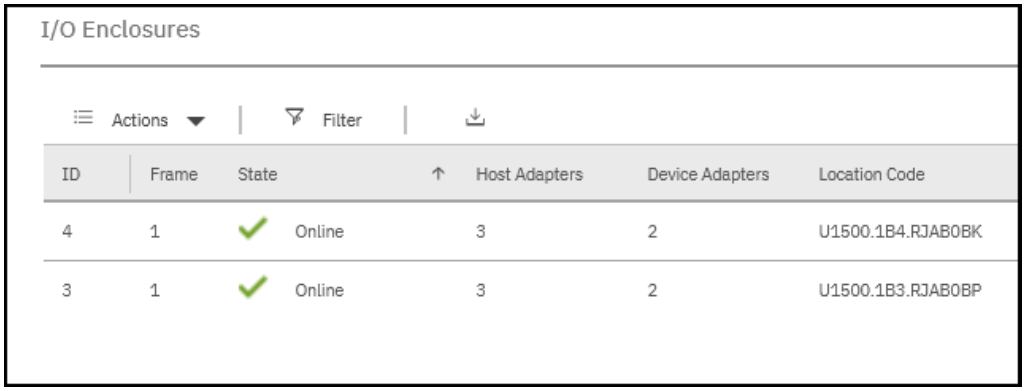


Figure 9-107 I/O Enclosures

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.
- Host adapter: Number of host adapters that are installed.
- DA: Number of DAs that are installed.

The *FC ports* are ports on a host adapter that connect the DS8A00 to hosts or another storage system either directly or through SAN switches/directors.

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 9-108.

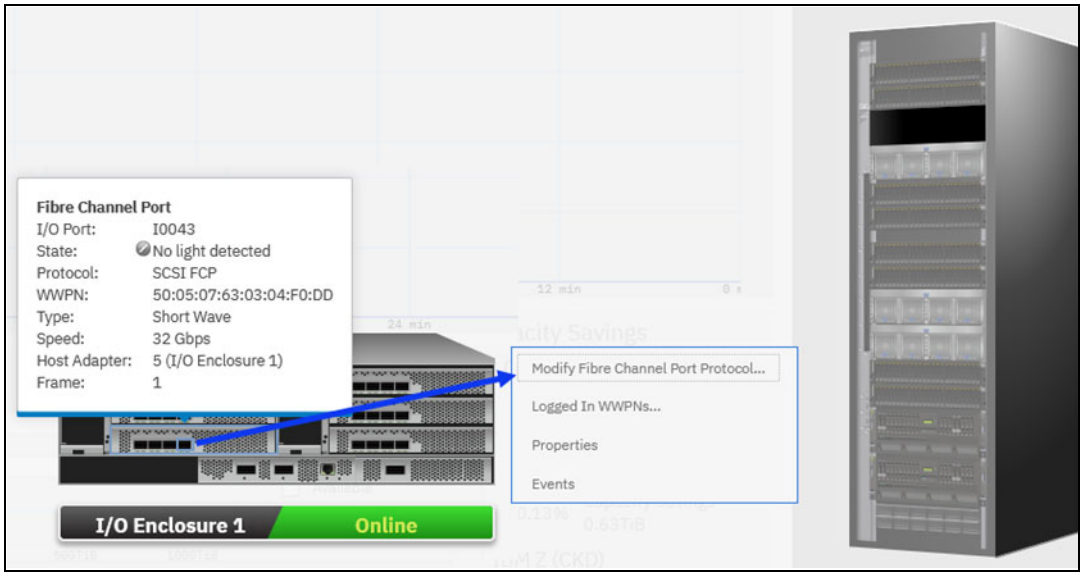


Figure 9-108 I/O enclosure showing host adapter and FC ports

The FC ports have several attributes:

- 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.

9.14.2 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 9-109 on page 314 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.

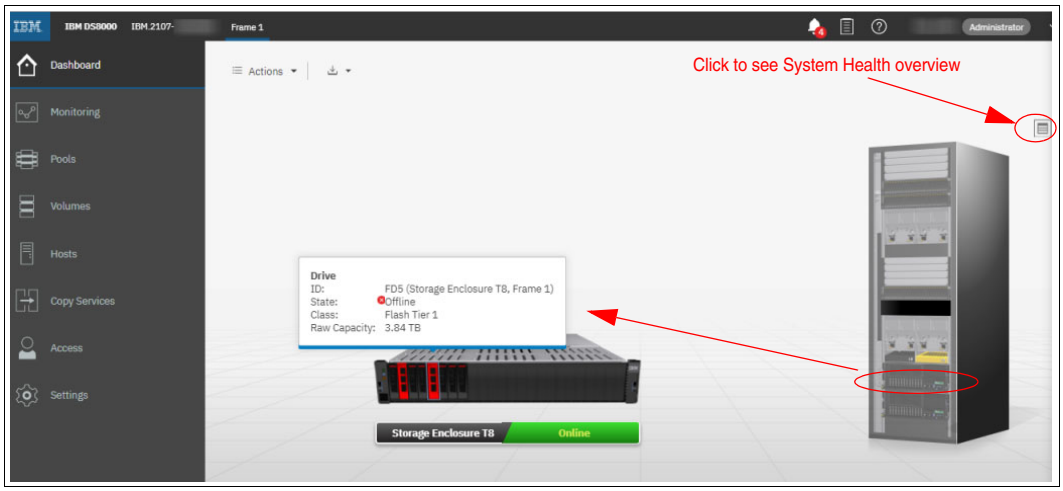


Figure 9-109 Hardware view with failed drives

The same information can be obtained by clicking the **Show System Health Overview** icon, as shown in Figure 9-109.

The System Health Overview opens and shows the failed state of the drives, as shown in Figure 9-110.

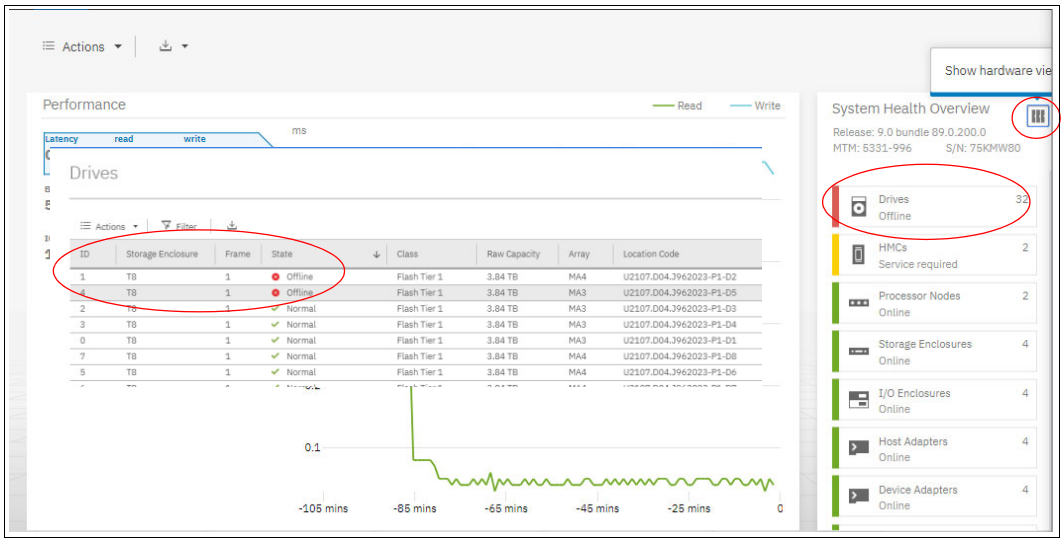


Figure 9-110 System Health Overview with failed drives

- Failed processor node: Figure 9-111 on page 315 shows the Hardware view with a failed processor node. Hovering your cursor over the node provides more information about this component.

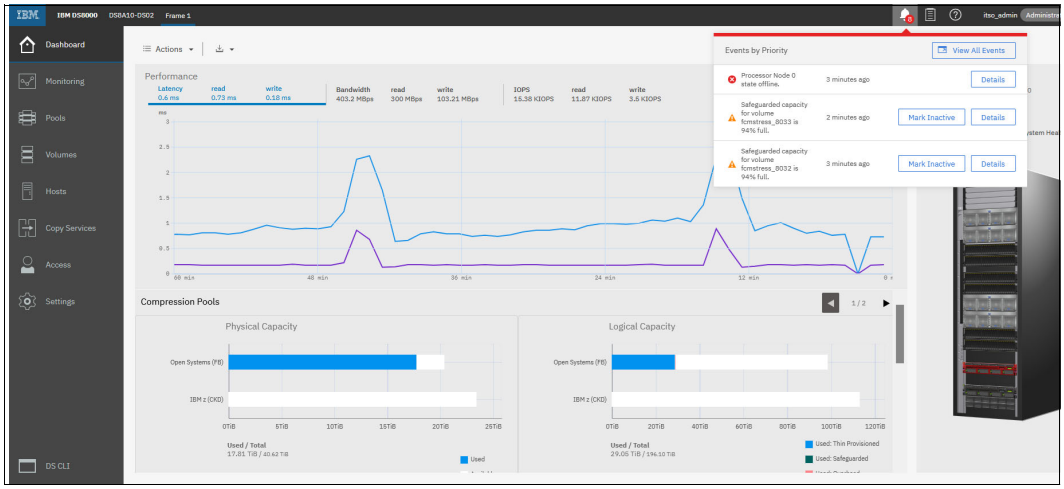


Figure 9-111 Hardware view with a failed processor node

The same information can be obtained from the System Health Overview, as shown in Figure 9-112.

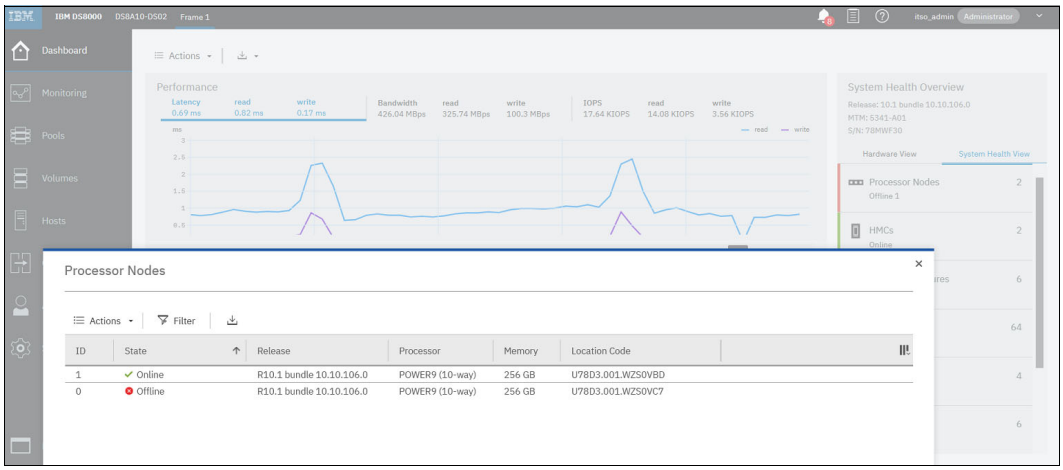


Figure 9-112 System Health Overview with a failed processor node

9.14.3 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 9-113.

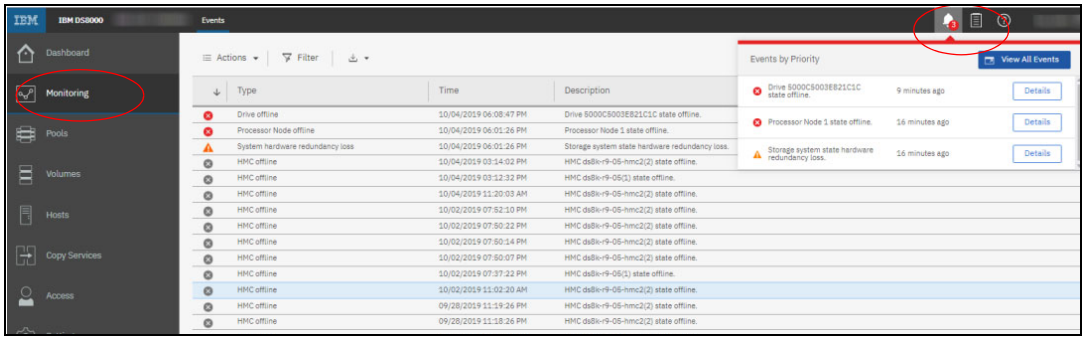


Figure 9-113 System Events window

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.

9.14.4 Exporting system-wide information

As shown in Figure 9-114, you can export detailed reports from the dashboard with the following system-wide information into a CSV file.

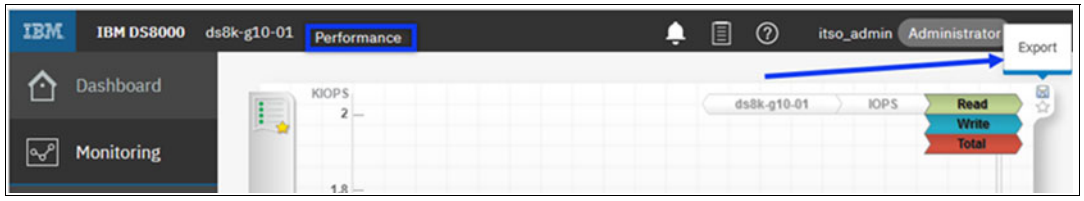


Figure 9-114 Export Storage System Performance sample

- 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 9-119 on page 322, and the DS CLI (see Section 10.4.8, “IBM Easy Tier” on page 388). 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*, 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 9-115 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 the 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 9-115 Example XLS format

9.14.5 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 DS8A00 can also be exported by selecting **Export Audit Log** on the Events window, as shown in Figure 9-116.

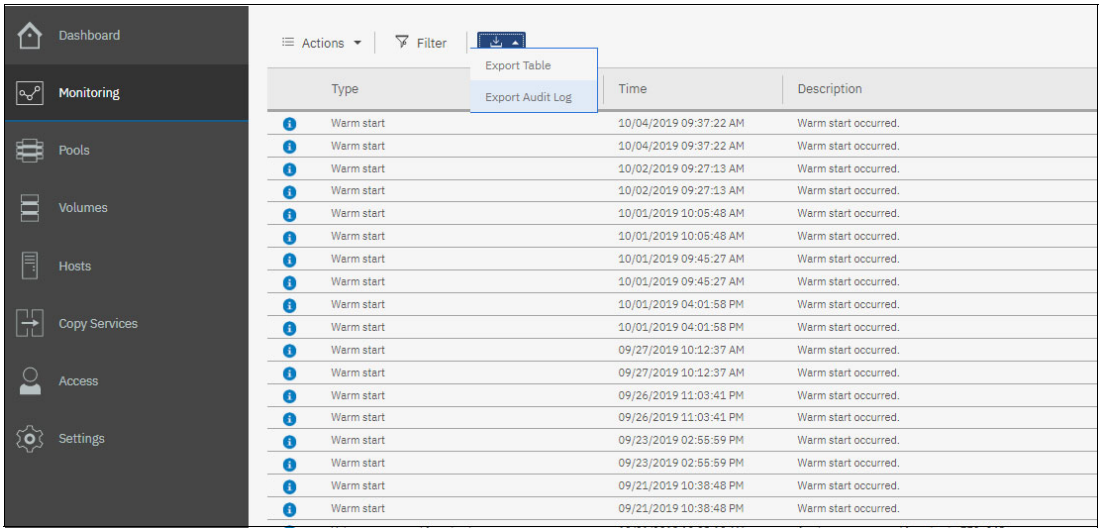


Figure 9-116 Export Audit Log

9.15 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 9-117. You can also set any of these graphs as the default.



Figure 9-117 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 9-118 shows a comprehensive view of the available performance functions.



Figure 9-118 Performance monitoring graph comprehensive view

To learn how to obtain statistics from the DS CLI, see Section 10.5, “Metrics with DS CLI” on page 388.

9.15.1 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 or IBM Storage Insights.

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 9-117 on page 319. You can 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 9-119.



Figure 9-119 StorageSystem_PerformanceSample export CSV file

9.15.2 Working with customized performance graphs

The DS GUI offers four predefined performance graphs, which are available in the Favorites menu, as shown in Figure 9-120.

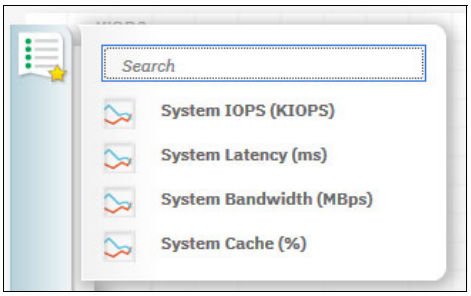


Figure 9-120 Predefined graphs

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 9-121.

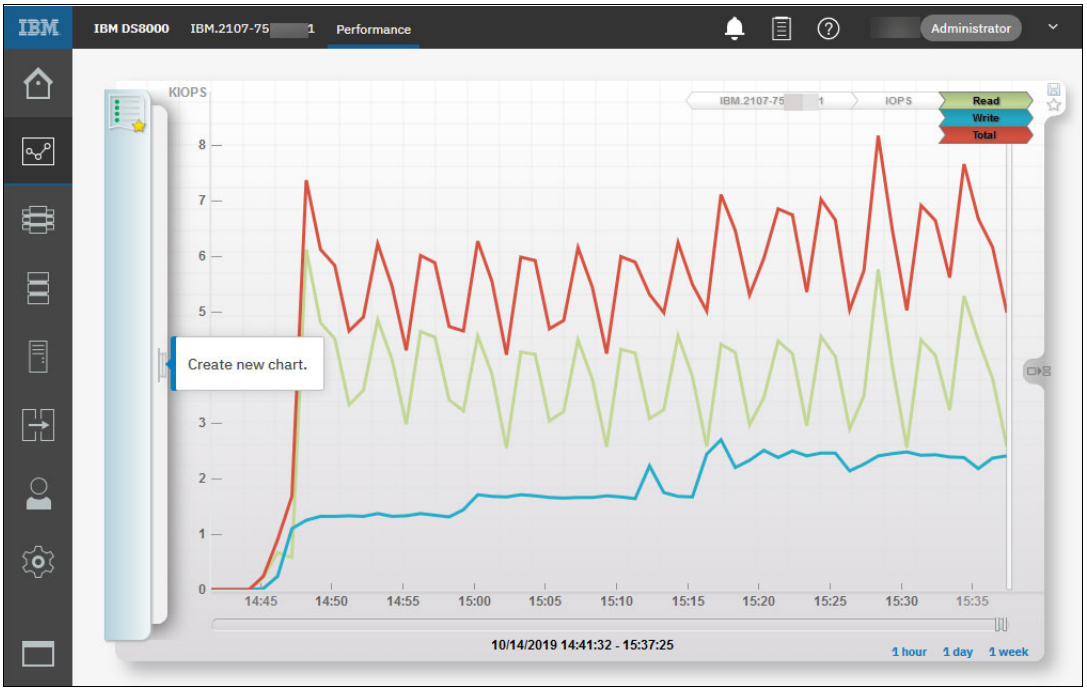


Figure 9-121 Creating a chart

Figure 9-122 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).

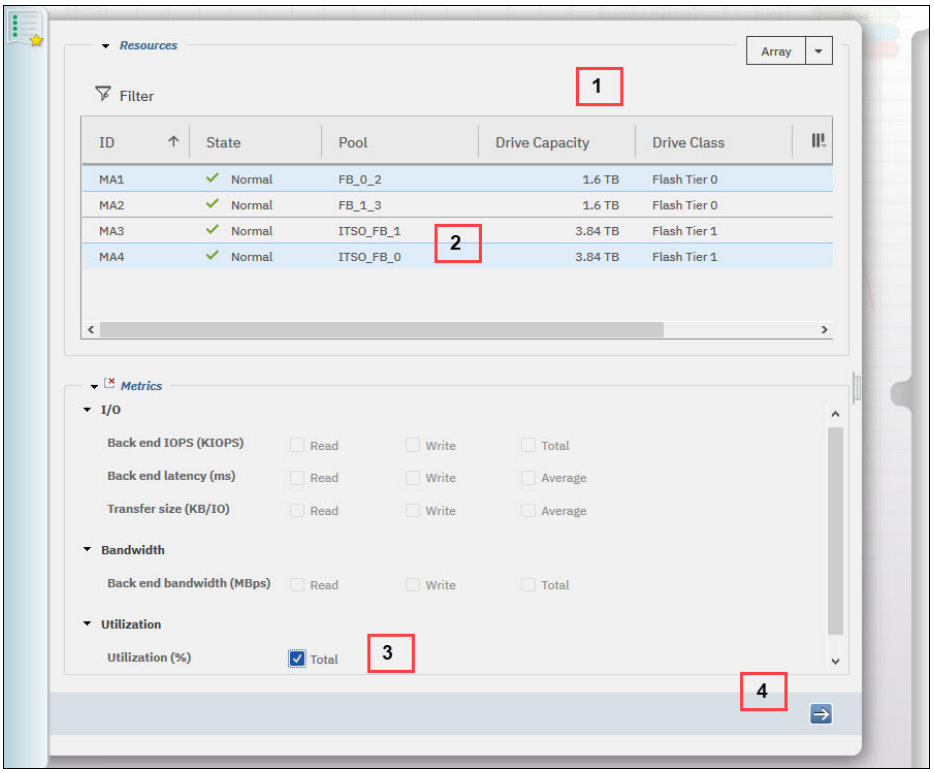


Figure 9-122 Array graph definition window

4. Draw the graph, which is shown in Figure 9-123.



Figure 9-123 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 9-121 on page 323, 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 9-124 shows the metric options that are available for the selected pool.

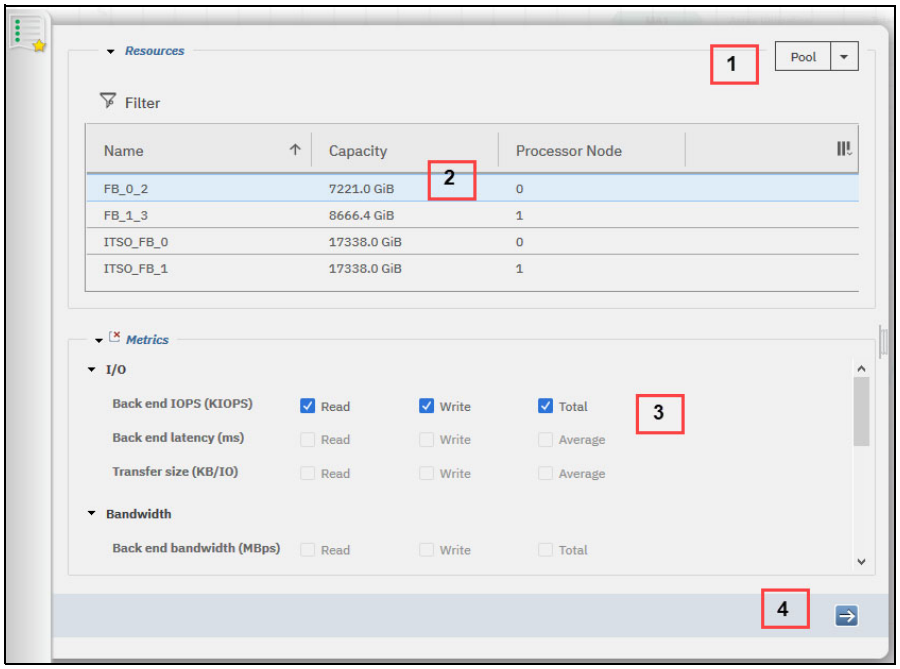


Figure 9-124 Pool graph definition window

4. Draw the graph, which is shown in Figure 9-125.

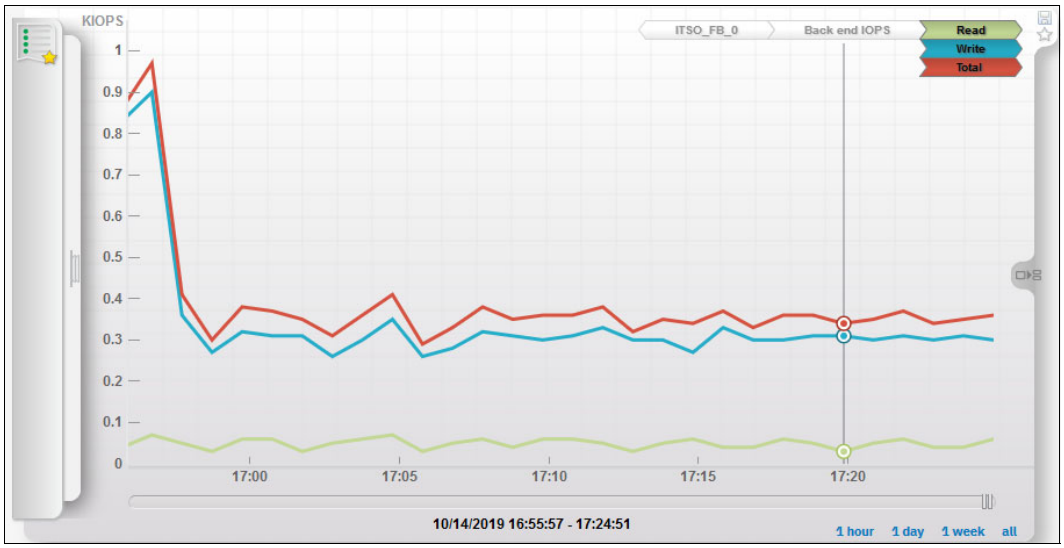


Figure 9-125 Pool performance graph

You can monitor Easy Tier directly in the DS GUI by using the workload categorization report and migration report. Figure 9-126 shows the Easy Tier pool level workload settings for creating total data movement report that is shown in Figure 9-127 on page 326.

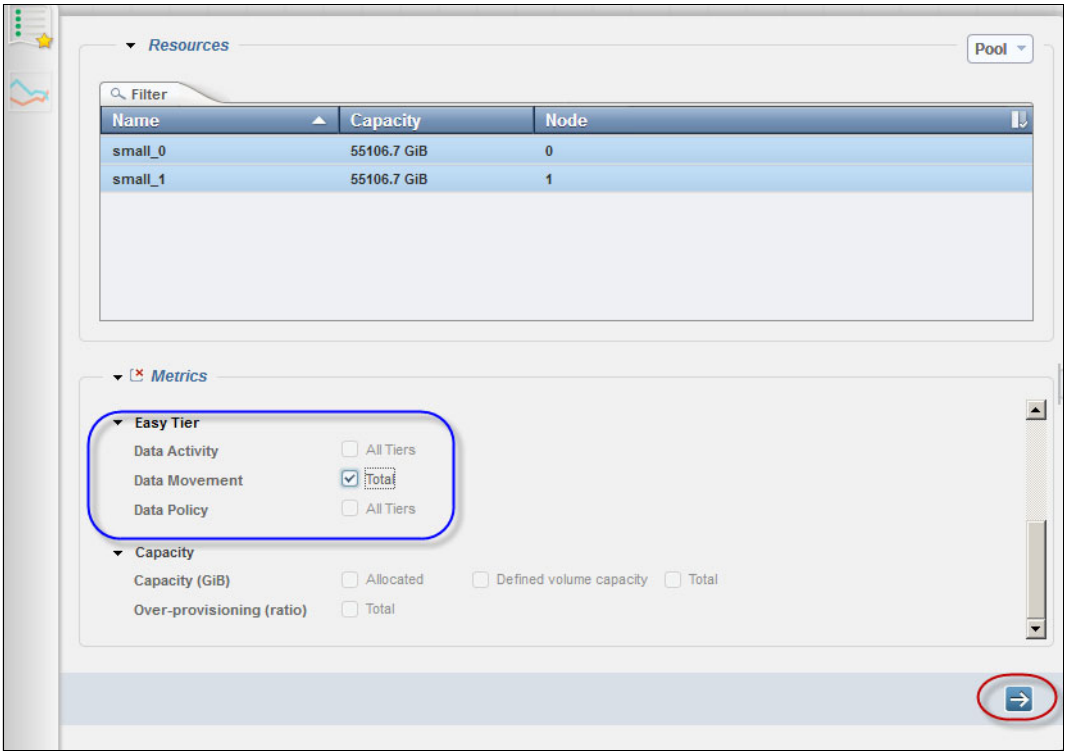


Figure 9-126 Example of Easy Tier settings for total Data Movement

Figure 9-127 shows an example of the Total Data Movement report.

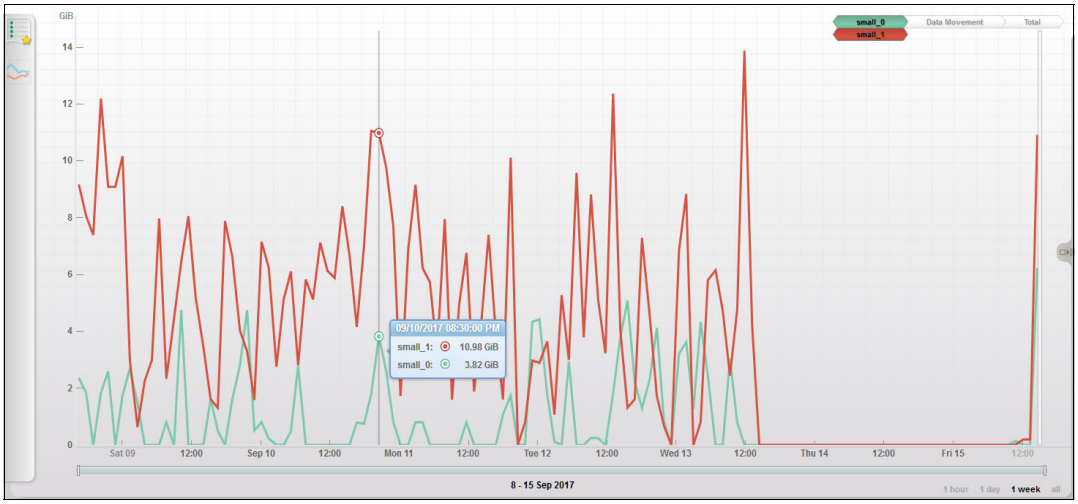


Figure 9-127 Easy Tier pool level workload reports (Total Data Movement)

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 9-128.

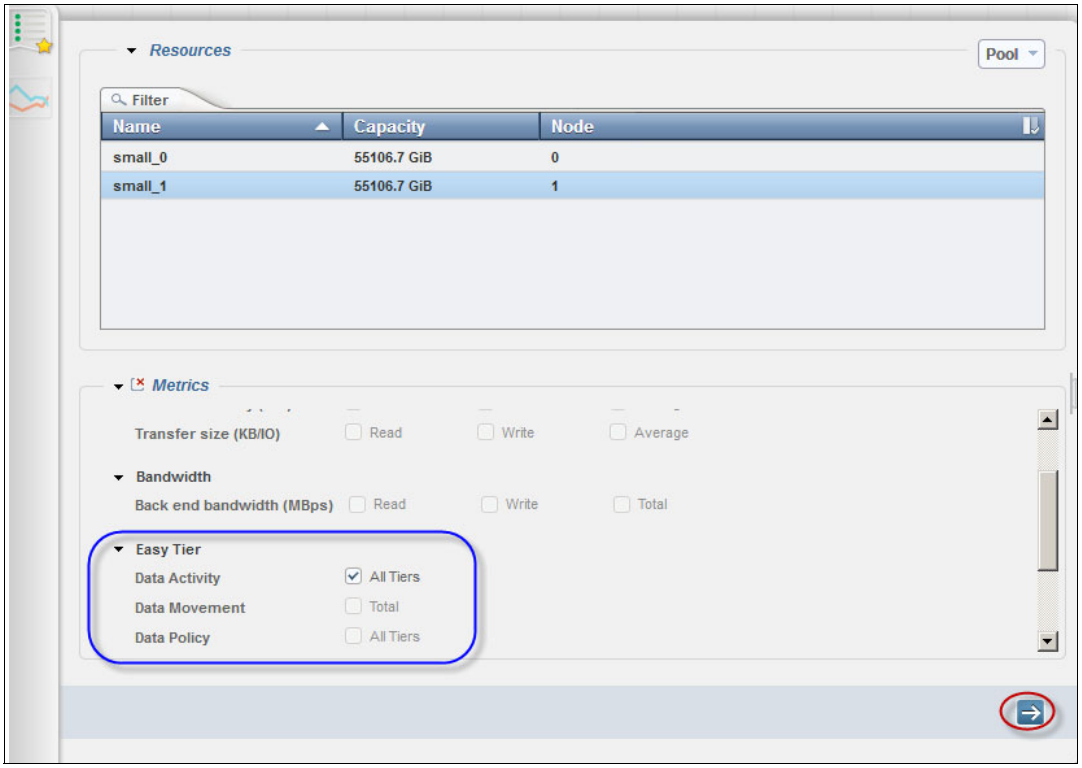


Figure 9-128 Example of Easy Tier Data Activity Report setup

An example report of Easy Tier Data Activity for a pool is shown in Figure 9-129.



Figure 9-129 Example of Easy Tier Data Activity Report

Creating system performance graphs

Figure 9-121 on page 323 shows the steps to start a chart. To create the graph of the system's performance as shown in Figure 9-130, 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.

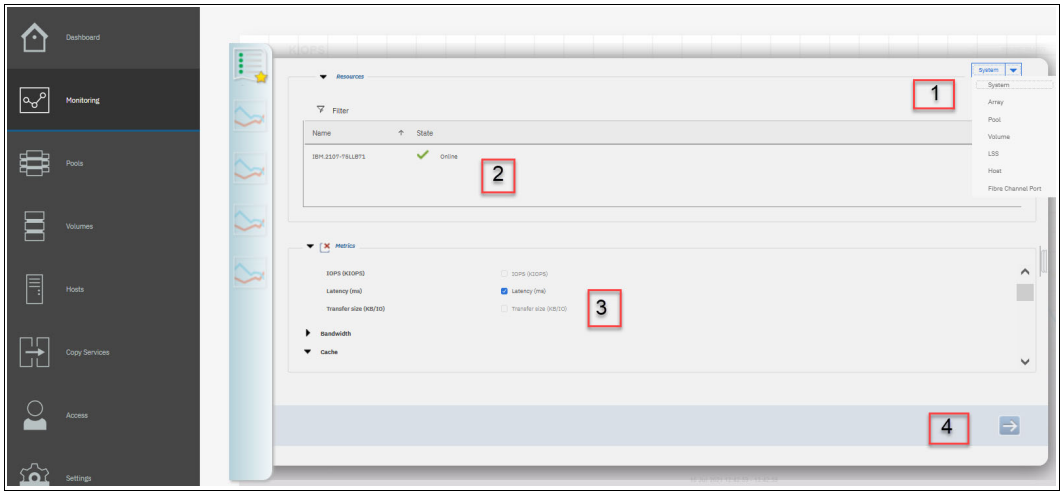


Figure 9-130 System graph definition window

4. Draw the graph, which is shown in Figure 9-131.

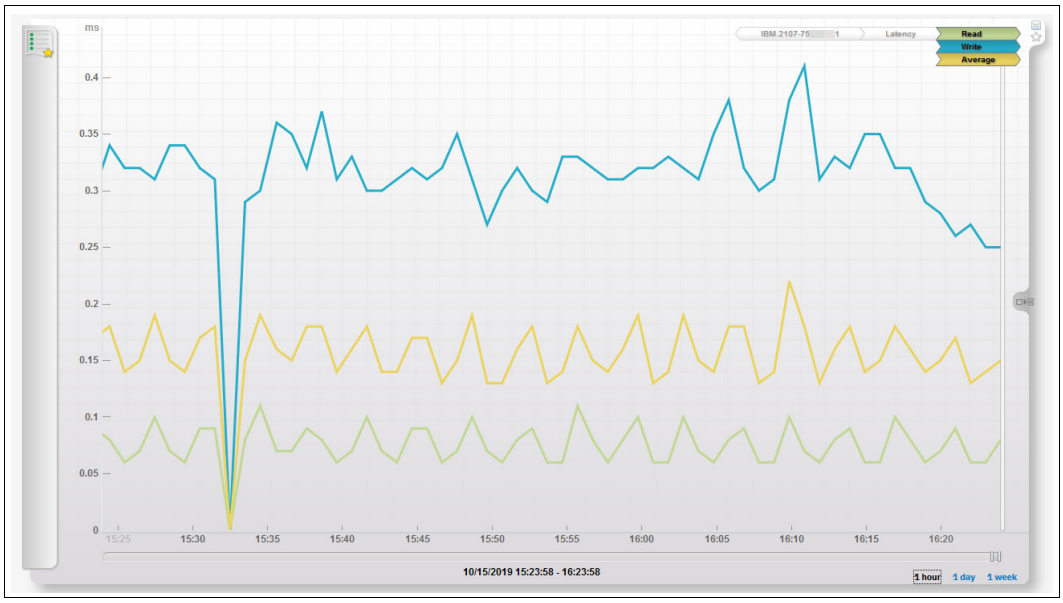


Figure 9-131 System performance graph

Creating FC port performance graphs

Start a chart as shown in Figure 9-121 on page 323. To complete the sequence that is indicated by the numbers in Figure 9-132 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.

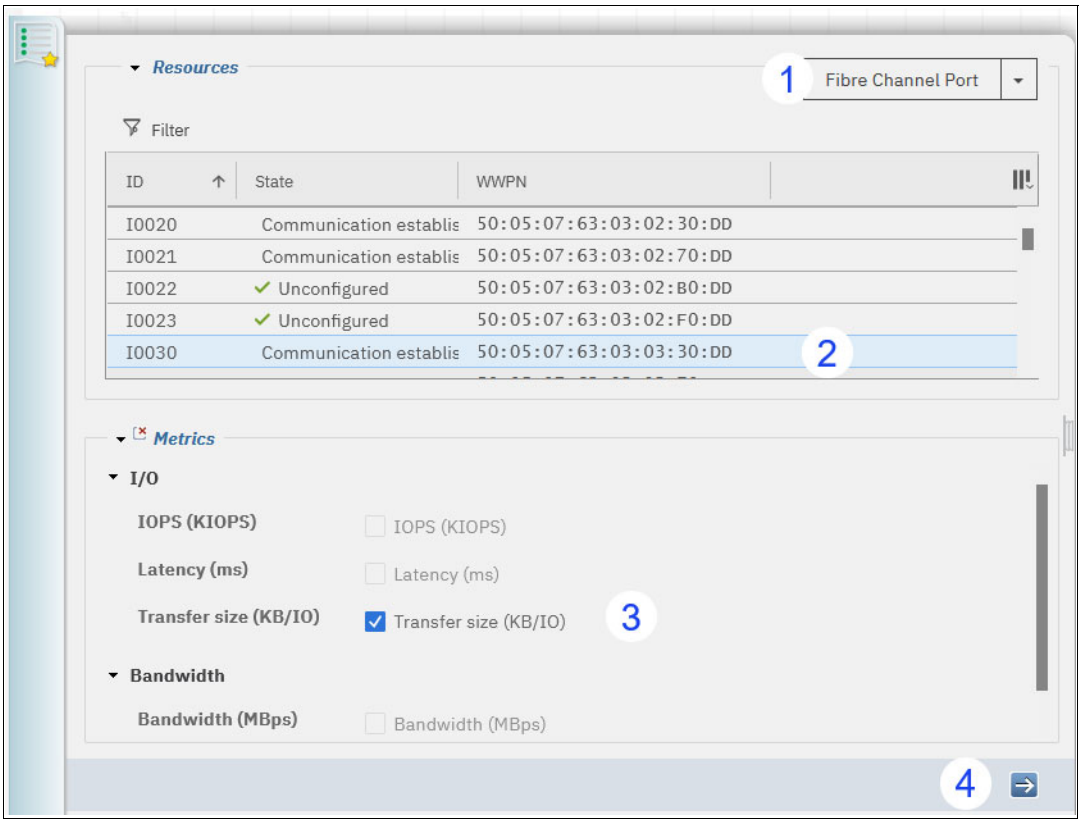


Figure 9-132 Fibre Channel Port graph definition window

4. Draw the graph. Figure 9-133 shows the completed graph.

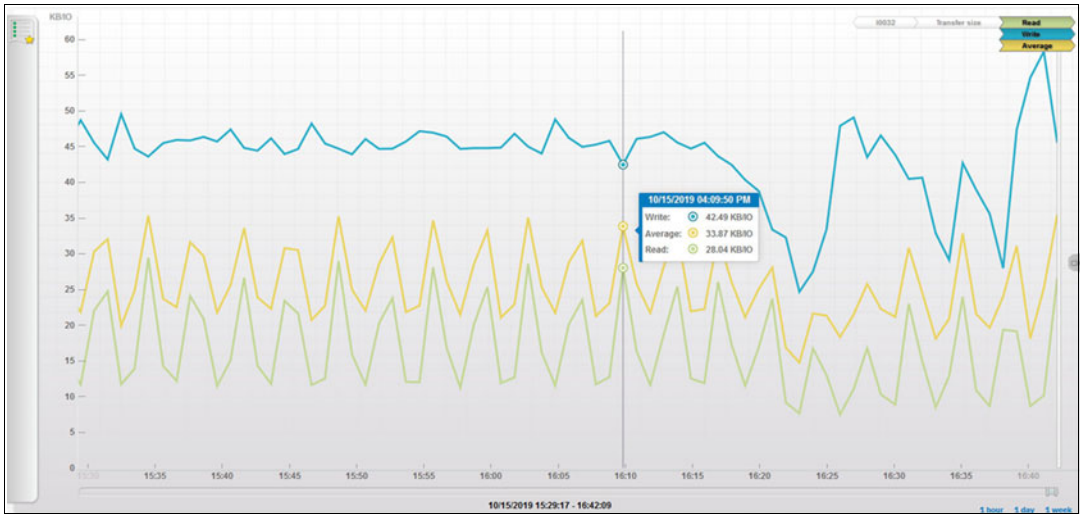


Figure 9-133 Fibre Channel Port performance graph

Creating volume performance graphs

Start a chart as shown in Figure 9-121 on page 323. To create the Easy Tier performance graphs for the volume, as shown in Figure 9-134, 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**.

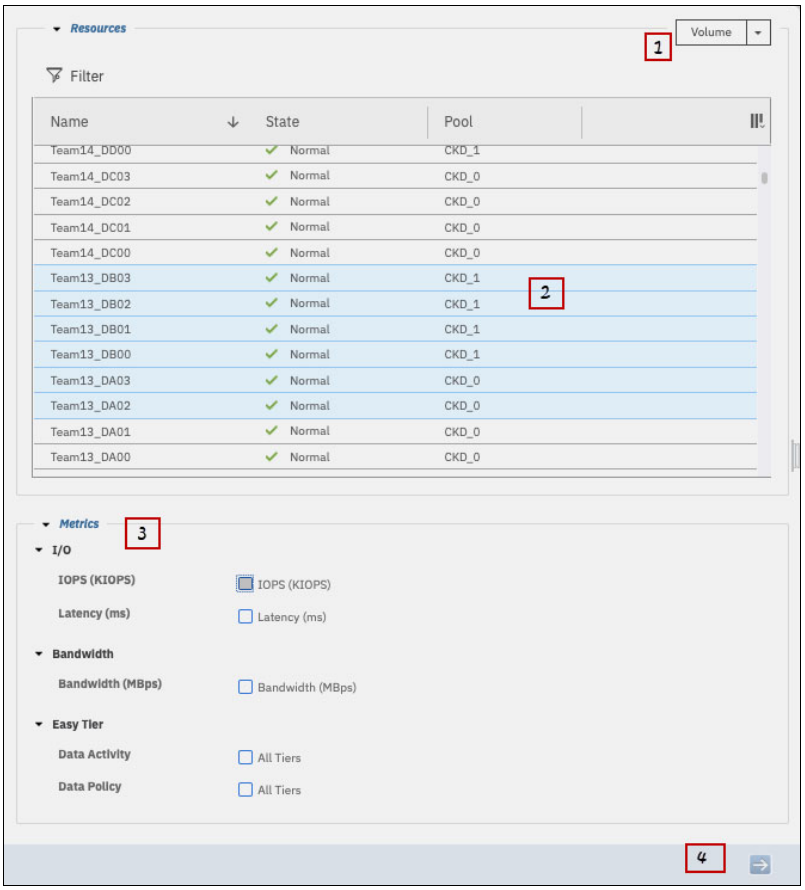


Figure 9-134 Volume Performance graph definition window

4. Draw the graph. Figure 9-135 shows the bandwidth graph for a set of five volumes.



Figure 9-135 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 9-136 shows the performance actions and metrics that are available for a volume on the Volume by LSS page.

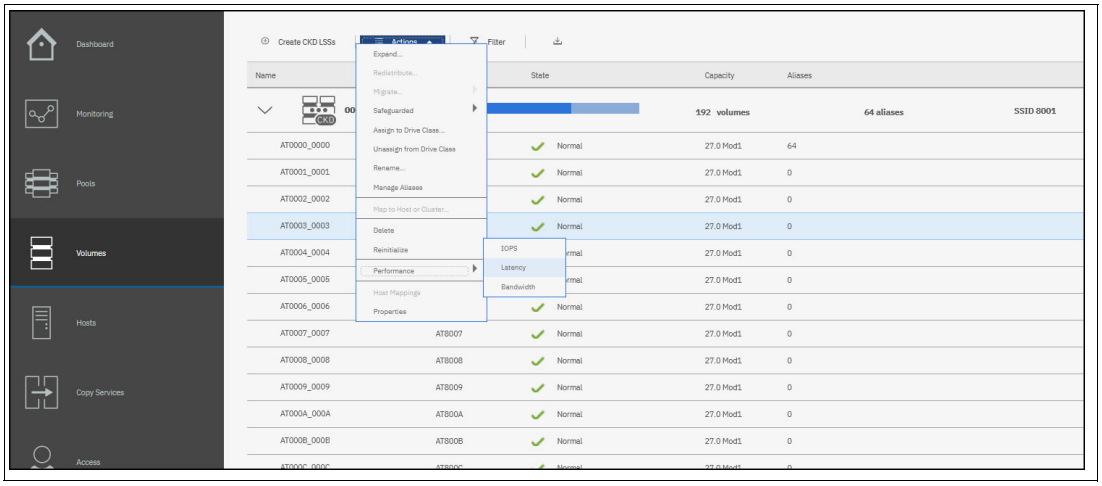


Figure 9-136 Performance action for a volume

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 9-137 on page 332:

1. Click the icon **Add to Favorites.**
2. Enter a name for it.
3. Search and display your favorite graphs from the list.

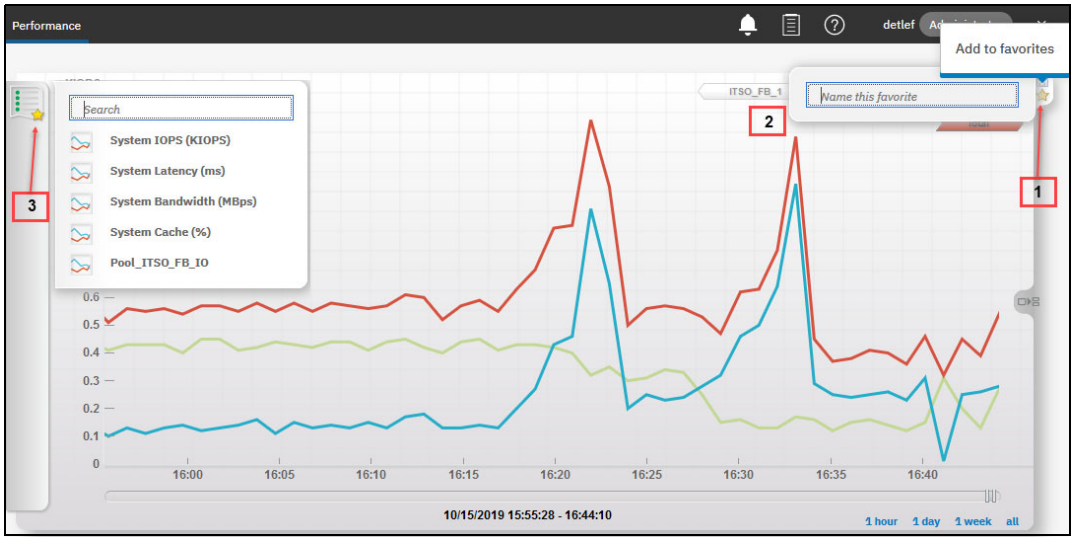


Figure 9-137 Adding a graph to the Favorites menu

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 9-138.

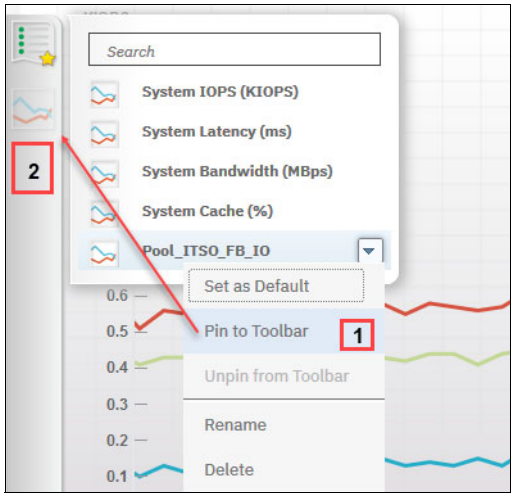


Figure 9-138 Pinning a graph to the toolbar

9.16 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 9-139. This feature is useful for a proactive check of SAN-related issues from a DS8A00 perspective.

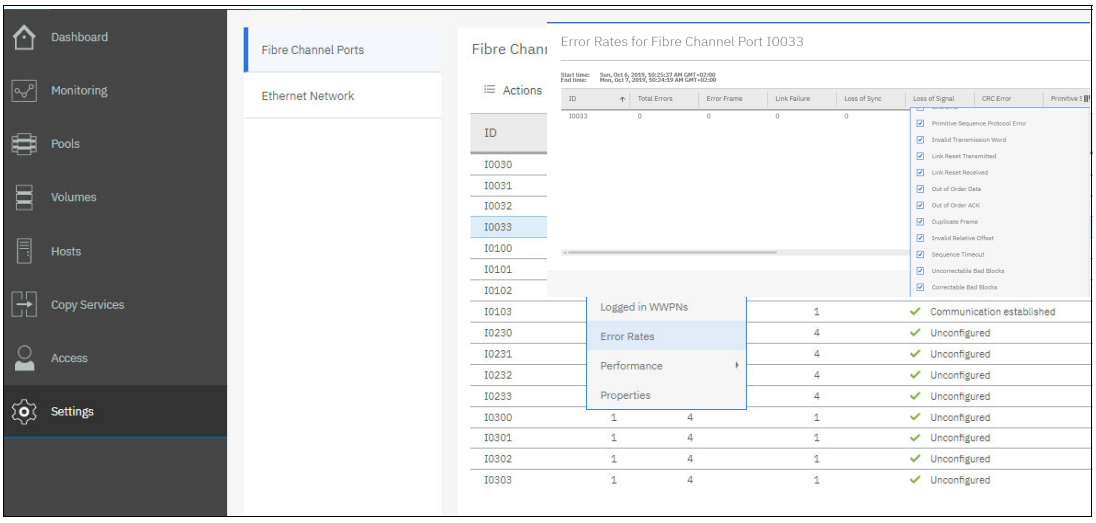


Figure 9-139 Fibre Channel Port error rate statistics

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 can also 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 can also 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 might also 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.

9.17 Providing feedback

Thirty days after the first login, a feedback request appears on the upper right under "Suggested tasks", as shown in Figure 9-140. 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 is snoozed for 90 days.

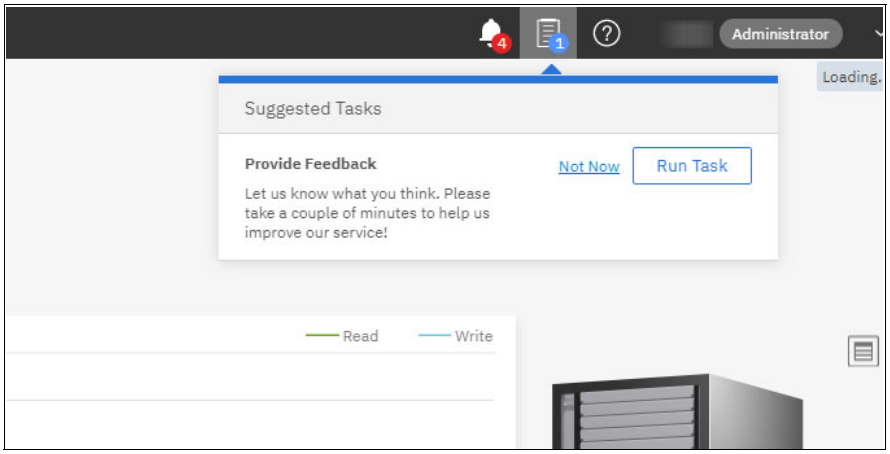


Figure 9-140 Suggested Tasks window

10

IBM DS8A00 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” on page 336
- “I/O port configuration” on page 353
- “DS8A00 storage configuration for Fixed-Block volumes” on page 354
- “DS8000 storage configuration for the CKD volumes” on page 375
- “Metrics with DS CLI” on page 388
- “Private network security commands” on page 396
- “Copy Services commands” on page 398
- “Earlier DS CLI commands and scripts” on page 398
- “Dual Control commands” on page 399
- “For more information” on page 402

Note: This chapter illustrates only a few essential commands. For a list of all commands and their parameters, see Command-line interface.

© Copyright IBM Corp. 2025. All rights reserved.

10.1 DS CLI overview

The DS CLI provides a full-function command set used for storage unit management and configuration of various application functions. For more information about DS CLI usage and setup, see Command-line interface.

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 installed on an external server and interacting with the DS8A00. Alternatively, the DS CLI embedded in the DS GUI can be used, and it can be accessed from the DS GUI dashboard by using the DS CLI option in the lower left corner 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 or Semeru. 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:

    ```
    dsclisetup.exe LAX_VM "C:\Program Files\java-whatever\jre\bin\java.exe"
    ```

    Note: Due to a space in the Program Files directory name, add quotation marks around the directory name.

  - For a UNIX or Linux system, specify the following path:

    ```
    dsclisetup.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 **dsclisetup.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=$LIBPATH unset LIBPATH
  dsclisetup.bin LAX_VM/opt/ibm-Java-whatever/java/bin/java
        export LIBPATH=$LIBSAVE
     unset LIBSAVE
  ```

10.1.3 Installing the DS CLI on a Windows 11 system

To install the DS CLI on a Microsoft Windows 11 system, make sure that the correct Java or Semeru version is installed. Example 10-1 shows how to install DS CLI on a Windows 11 system with the extra step of passing the JVM path location. You can download the DS CLI from IBM Fix Central and download the IBM Semeru runtime from IBM Semeru Runtime.

Example 10-1 DS CLI installation with a path to Semeru on a Windows 11 system

```
C:\IMAGES\HMC\Disk1\InstData\Windows\NoVM>dsclisetup.exe LAX_VM "C:\Program
Files\Semeru\jre-8.0-openj9\bin"
```

For instances where Semeru or Java is already set up, the installation starts with running the **dsclisetup.exe** program that is found in your installation media, as shown in Figure 10-1.

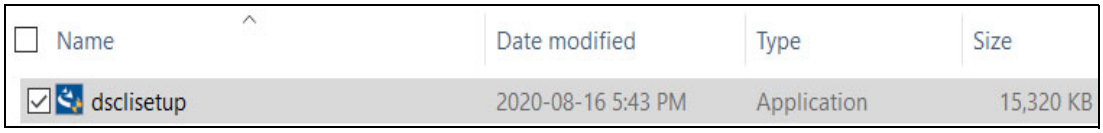


Figure 10-1 The dsclisetup.exe file

For more information about IBM Semeru, see IBM Semeru Runtime.

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 HIWN63A. 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*, GI13-3563. You can use the order number (GI13-3563) 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` and 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/CSMDSCLI`. 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/CSMDSCLI
```

```
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/CSMDSCLI

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/CSMDSCLI 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=TSO       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=TSO       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> lssi
Name          ID           Storage Unit    Model WWNN          State  ESSNet
==============================================================================================
IBM.2107-75xxxx IBM.2107-75xxxx IBM.2107-75xxxxx A50 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.2107-75xxxxx -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? Ý 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/CSMDSCLI/dscli -cfg /opt/IBM/CSMDSCLI/profile/aca91.profile
    ver -l;
echo "Command 2:";
/opt/IBM/CSMDSCLI/dscli -cfg /opt/IBM/CSMDSCLI/profile/aca91.profile
    lsrank;
echo "Command 3:";
/opt/IBM/CSMDSCLI/dscli -cfg /opt/IBM/CSMDSCLI/profile/aca91.profile
    lsarray;
/*
```

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 **-l** parameters together.

-l (optional): The **-l** parameter displays the versions of the DS CLI client program, Storage Manager, HMC code level, LMC, and the code bundle. You cannot use the **-l** 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

-stgmgr (optional): The **-stgmgr** 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 an earlier DS CLI version.

Example 10-4 DS CLI version command

```
dscli> ver -l
Date/Time: 30 October 2025 18:38:45 CET IBM DSCLI Version: 7.10.10.348 DS: -
DSCLI          7.10.10.348
StorageManager 10.2.2507181727
HMC DSCLI      7.10.10.348
================Version=================
Storage Image    LMC         Bundle Version
=========================================
IBM.2107-78HAL91 7.10.10.348
```

The following DS CLI commands check the IBM Copy Services Manager version and install new software levels of IBM Copy Services Manager:

- The **lssoftware -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> lssoftware -l
Date/Time: 30 October 2025 18:41:10 CET IBM DSCLI Version: 7.10.10.348 DS:
IBM.2107-78HAL91
Type Version                  Status  HMC
=======================================
 CSM V6.3.15.0-a20250818-1146 Running 1
 CSM V6.3.15.0-a20250818-1146 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-6.3.15.0-linux-ppcle.bin -certloc
/home/hscroot/wht/csm-setup-6.3.15.0-linux-ppcle.bin.crt -type CSM -hmc 1
IBM DSCLI Version: 7.10.10.348 DS: IBM.2107-78xxxxx
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 DS8A00 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 6.6, “User management” on page 184.

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 DS8A00 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 `John`, which is in the `op_storage` group. The temporary password of the user is `passw0rd`. The user must change the temporary password during the first login.

Example 10-7 Using the mkuser command to create a user

```
dscli> mkuser -pw passw0rd -group op_storage John
```

```
CMUC00133I mkuser: User John successfully created.
```

- **rmuser**

  An existing user ID is removed by using this command. Example 10-8 shows the removal of a user called Jane.

  Example 10-8 Removing a user

```
dscli> rmuser Jane
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 can also 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 John
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 -l
Name                Group           State  Scope
================================================
admin               admin           active *
engineering         ibm_engineering active *
John                op_storage      active PUBLIC
secadmin            secadmin        active *
service             ibm_service     active *
suser               ibm_service     active PUBLIC
```

- **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 John
Name         John
Group        op_storage
State        active
FailedLogin  0
DaysToExpire 9999
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 John -pw passw0rd
CMUC00206I managepwfile: Record 10.0.0.1/John 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        8
Password History      8
```

10.1.8 DS CLI profile

To access the DS8A00 with the DS CLI, you must provide certain information by using the **dscli** command. At a minimum, the IP address or hostname of the DS8A00 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 storeage_image_ID" command
options, respectively.
#devid:  IBM.2107-AZ12341
#remotedevid:IBM.2107-AZ12341

devid:    IBM.2107-75xxxxx
hmc1:     10.0.0.1
username: admin
pwfile:   c:\mydir\75xxxxx\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 **dscli** commands that require it. Additionally, if you specify the `dscli` 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/dscli/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 **dscli** output:

- **banner**: Date and time with the **dscli** 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\dscli\profile` to `C:\Program Files (x86)\IBM\dscli\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\dscli>dscli -hmc1 10.0.0.1 -hmc2 10.0.0.5
Enter your user name: John
Enter your password: xxxxx
DS: IBM.2107-75xxxxx
dscli>
```

Alternatively, you can modify the following lines in the `dscli.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 DS CLI commands have a common structure which can include up to four components arranged in the following order:

1. The **command name** specifies the task that the DS CLI performs. For example, **`lsarray`** lists all defined arrays in the storage system.
2. **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. They are always preceded by a hyphen (-) and can be followed immediately by a space and a flag parameter. For example, **`-h`** provides additional help with command syntax.
3. **Flags parameter** provides information that is required to implement the command modification that is specified by a flag. They are variables as their value changes based on your specific need. For example, flag **`-user`** requires a flag parameter that is a specific *user_name*
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 it is not preceded by a flag.

In Example 10-18 on page 350, **`lsrank`** is the command name, **`-dev`** and **`-l`** are command parameters and *IBM.2107–75xxxxx* is the sub parameter for the **`-dev`** parameter, and *R1*, *R2*, and *R3* are a list of command parameters.

Example 10-18 DS CLI components and structure

```
dscli> lsrank -dev IBM.2107-75xxxxx -l R1 R2 R3
```

10.1.11 Using the DS CLI application

To issue commands to the DS8A00, you must first log in to the DS8A00 through the DS CLI client 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 can 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:
   - **dscli -hmc1 <hostname or ip address> -user <adm user> -passwd <pwd> <command>**
   - **dscli -cfg <dscli profile> -pwfile <security file> <command>**

   Important: For security reasons, 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 10-19 shows the use of the single-shot command mode.

Example 10-19 Single-shot command mode

```
C:\Program Files (x86)\IBM\dscli>dscli -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 DS8A00 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 **dscli** command prompt.
3. Use the DS CLI commands and parameters. You are not required to begin each command with **dscli** because this prefix is provided by the **dscli** 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 DS CLI profile file by using the **paging:off** parameter.

Example 10-20 shows the interactive command mode by using the profile `DS8A00.profile`.

Example 10-20 Interactive command mode

```
C:\Program Files (x86)\IBM\dscli>dscli -cfg DS8A00.profile

dscli> lsarraysite -l
Date/Time: 20 September 2024 5:42:24 PM IBM DSCLI Version: 7.10.0.767 DS: IBM.2107-78xxxxx
arsite DA Pair dkcap (10^9B) diskrpm State      Array diskclass       encrypt
================================================================================
S1     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S2     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S3     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S4     0              4800.0   65000 Unassigned -     FlashCoreModule supported

dscli> lssi
Name        ID               Storage Unit     Model WWNN             State  ESSNet
==============================================================================
ds8k-g10-01 IBM.2107-78xxxxx IBM.2107-78xxxxx A05   5005076303FFF0DD Online Enabled
```

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 10-21 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, as well as in the embedded DS CLI in the DS GUI. Only one authentication process is needed to run all of the script commands.

Example 10-21 Example of a DS CLI script file

```
# Sample dscli script file
# Comments can appear if hashed
lsarraysite
lsarray
lsrank
```

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 10-22 shows starting the DS CLI by using the **-script** parameter and specifying a profile and script path and name that contains the commands from Example 10-21 on page 351.

Example 10-22 Running a DS CLI file

```
C:\Program Files (x86)\IBM\dscli>dscli -cfg DS8A00.profile -script c:\ds8000.script
arsite DA Pair dkcap (10^9B) diskrpm State      Array diskclass       encrypt
================================================================================
S1     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S2     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S3     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S4     0              4800.0   65000 Unassigned -     FlashCoreModule supported
Array State      Data   RAIDtype    arsite Rank DA Pair DDMcap (10^9B) diskclass       encrypt
=============================================================================================
A0    Unassigned Normal 6 (5+P+Q+S) S1     -    0               4800.0 FlashCoreModule supported
A1    Unassigned Normal 6 (5+P+Q+S) S2     -    0               4800.0 FlashCoreModule supported
ID Group State      datastate Array RAIDtype extpoolID stgtype
=============================================================
R0     - Unassigned Normal    A0           6 -         fb
R1     - Unassigned Normal    A1           6 -         fb
```

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/messages that are used by the DS CLI are listed in Messages.

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 10-23 lists the I/O ports by using the **lsioport** command. I0020–I0023 are on one adapter, and I0030–I0033 are on another adapter.

Example 10-23 Listing the I/O ports

```
dscli> lsioport
ID    WWPN             State                      Type             topo     portgrp Security
==========================================================================================
I0020 50050763030230DD Offline                    Fibre Channel-LW -        0       Disabled
I0021 50050763030270DD Offline                    Fibre Channel-LW -        0       Disabled
I0022 500507630302B0DD Offline                    Fibre Channel-LW -        0       Disabled
I0023 500507630302F0DD Offline                    Fibre Channel-LW -        0       Disabled
I0030 50050763030330DD Offline                    Fibre Channel-LW -        0       Disabled
I0031 50050763030370DD Offline                    Fibre Channel-LW -        0       Disabled
I0032 500507630303B0DD Offline                    Fibre Channel-LW -        0       Disabled
I0033 500507630303F0DD Offline                    Fibre Channel-LW -        0       Disabled
I0040 50050763030430DD Offline                    Fibre Channel-SW -        0       Disabled
I0041 50050763030470DD Offline                    Fibre Channel-SW -        0       Disabled
I0042 500507630304B0DD Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
I0043 500507630304F0DD No light detected          Fibre Channel-SW SCSI-FCP 0       Disabled
I0050 50050763030530DD Offline                    Fibre Channel-SW -        0       Disabled
I0051 50050763030570DD Offline                    Fibre Channel-SW -        0       Disabled
I0052 500507630305B0DD Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
I0053 500507630305F0DD No light detected          Fibre Channel-SW SCSI-FCP 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.

The *Security* field in the **lsioport** output (Example 10-23) indicates the status of the IBM Fibre Channel Endpoint Security feature (In-flight data encryption over FICON links). For more information, see *IBM Fibre Channel Endpoint Security for IBM DS8000 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 10-24 shows setting I/O ports without and with the **-force** option to the FICON topology, and then checking the results.

Example 10-24 Changing the topology by using setioport

```
dscli> lsioport i0103
ID    WWPN             State                      Type             topo     portgrp Security
==========================================================================================
```

```
I0103 500507630908D462 Communication established Fibre Channel-SW SCSI-FCP 0       Disabled
dscli> 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.
dscli> 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 388.

10.3 DS8A00 storage configuration for Fixed-Block volumes

This section reviews examples of a typical DS8A00 storage configuration when the DS8A00 storage is attached to open system hosts. You can perform the DS8A00 storage configuration by completing the following steps:

1. Create the arrays.
2. Create the ranks.
3. Create the extent pools.
4. Create the volumes.
5. Create the volume groups.
6. Create the host connections.
7. 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 DS8A00:

- The FlashTier 1 flash disk class specifies 2.5-inch 1.6 TB flash drives.
- The FlashTier 2 flash disk class specifies 2.5-inch 1.92 TB, 3.84 TB, or 7.68 TB flash drives.
- The FlashCoreModule flash disk class specifies 2.5-inch compression capable flash drives with sizes of 4.8 TB, 9.6 TB, or 19.2 TB.

Important: For more information about the current drive choices and RAID capacities, see IBM Documentation.

10.3.2 Creating the arrays

This section describes steps for creating the arrays. *Array sites* are groups of eight drives that are predefined in the DS8A00.

Important: 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.

The **lsarraysite** command displays all available array sites that are installed in the DS8A00. In Example 10-25, all listed array sites are unassigned. Therefore, all array sites can be used for creating arrays. The **-l** option reports the `diskclass` information.

Example 10-25 Listing array sites

```
dscli> lsarraysite -l
arsite DA Pair dkcap (10^9B) diskrpm State      Array diskclass       encrypt
=============================================================================
S1     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S2     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S3     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S4     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S5     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S6     0              4800.0   65000 Unassigned -     FlashCoreModule supported
S7     1              1600.0   65000 Unassigned -     FlashTier1      supported
S8     1              1600.0   65000 Unassigned -     FlashTier1      supported
```

You can issue the **mkarray** command to create arrays, as shown in Example 10-26. 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 10-26 Creating arrays by using mkarray

```
dscli> mkarray -raidtype 6 -arsite S1
CMUC00004I mkarray: Array A0 successfully created.
dscli> mkarray -raidtype 6 -arsite S2
CMUC00004I mkarray: Array A1 successfully created.
```

You can now see the arrays that were created by using the **lsarray -l** command, as shown in Example 10-27.

Example 10-27 Listing the arrays by using lsarray

```
dscli> lsarray -l
Array State      Data   RAIDtype    arsite Rank DA Pair DDMcap (10^9B) diskclass       encrypt
=============================================================================================
A0    Unassigned Normal 6 (5+P+Q+S) S1     -    0              4800.0 FlashCoreModule supported
A1    Unassigned Normal 6 (5+P+Q+S) S2     -    0              4800.0 FlashCoreModule supported
```

Example 10-28 shows the results of the **lsarraysite -l** command.

Example 10-28 Listing the high-performance flash disk class flash Tier 0 by using the lsarraysite command

```
dscli> lsarraysite -l -diskclass FlashCoreModule
arsite DA Pair dkcap (10^9B) diskrpm State      Array diskclass       encrypt
=============================================================================
S1     0              4800.0   65000 Assigned   A0    FlashCoreModule supported
S2     0              4800.0   65000 Assigned   A1    FlashCoreModule supported
```

You can see the following information in the previous examples:

- 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 (4.8 TB).
- Array sites (S1 and S2) that were used to create the arrays.
- Disk class (FlashCoreModule).

RAID types

RAID 6 and RAID 10 are supported RAID types for DS8A00 systems. RAID 5 is not supported with DS8A00. You receive an alert and error message if you try to create RAID 5. Example 10-29 shows an attempt to create RAID 5.

Example 10-29 Alert and error message when attempting to create RAID 5 array

```
dscli> mkarray -raidtype 5 -arsite S3

CMUC00536W mkarray: The use of RAID 6 over RAID 5 is highly recommended for
increased reliability. Please acknowledge that you understand the risks associated
with RAID 5.: Are you sure you want to accept the disclaimer above? [Y/N]: y
CMUN81522E mkarray: RAID 5 is not supported by the device adapter pair.
```

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 for creating Fix Block - FB rank is **mkrank -array A*x* -stgtype FB**.

If you enabled Data-at-Rest encryption with local or external key manager, then when you create ranks, you must include the **-keygrp** flag in the **mkrank** command with the corresponding *Encryption Key Group* **ID** parameter which is included in the **lskeygrp** command output.

After all the ranks are created, use the **lsrank** command to display the following information:

- List of all created ranks with associated ID.
- The rank group 0 or 1. Each group has server affinity (0 for server0 or 1 for server1). Rank group ID is displayed only when a rank is assigned to a specific extent pool.
- The RAID type.
- The format of the rank (fb or ckd).

Example 10-30 shows the **lskeygrp** and **mkrank** commands and the result of a successful **lsrank** command.

Example 10-30 Creating and listing ranks by using the mkrank and lsrank commands

```
dscli> lskeygrp
ID  state      reckeystate reckeydate datakeydate keyprotocol type name
==========================================================================
1   accessible disabled    -          09/18/2024  LOCAL       DAR  DAR_1

As per the above output, the Encryption Key Group ID = 1. It must be specified in
the mkrank command in -keygrp parameter.

dscli> mkrank -array A0 -keygrp 1 -stgtype fb
CMUC00007I mkrank: Rank R0 successfully created.
dscli> mkrank -array A1 -keygrp 1 -stgtype fb
CMUC00007I mkrank: Rank R1 successfully created.
```

```
dscli> lsrank
Date/Time: 20 September 2024 6:27:18 PM IBM DSCLI Version: 7.10.0.767 DS: IBM.2107-78xxxxx
ID Group State      datastate Array RAIDtype extpoolID stgtype
==========================================================
R0     - Unassigned Normal    A0           6 -         fb
R1     - Unassigned Normal    A1           6 -         fb

dscli> lsarray
Array State    Data   RAIDtype    arsite Rank DA Pair DDMcap (10^9B)
=================================================================
A0    Assigned Normal 6 (5+P+Q+S) S1     R0   0               4800.0
A1    Assigned Normal 6 (5+P+Q+S) S2     R1   0               4800.0
```

When defining a rank, you can also specify the extent size. You can have ranks and extent pools with large 1 (gibibyte) 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.

Small extents are the default.

10.3.4 Creating the extent pools and assigning ranks to 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 specified by the **-stgtype** parameter. **FB** is used for Fixed-Block format.
- 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 a best practice to note the ID.

The **mkextpool** command forces you to name the extent pools. To do so, complete the following steps (Example 10-31 on page 358):

1. Create empty extent pools by using the **mkextpool** command. The **keygrp** parameter must be used because encryption is enabled. See Example 10-30 on page 356 for more information in finding the encryption group ID value.
2. List the extent pools to obtain their IDs using **lsextpool** command.
3. Assign a rank to an extent pool by using the **chrank** command.

4. List the extent pools again with **lsextpool**. Note the capacity changes after assigning ranks.

Example 10-31 Creating an extent pool by using mkextpool, lsextpool, and chrank

```
dscli> mkextpool -rankgrp 0 -keygrp 1 -stgtype fb FB_0
CMUC00000I mkextpool: Extent pool P0 successfully created.
dscli> mkextpool -rankgrp 1 -keygrp 1 -stgtype fb FB_1
CMUC00000I mkextpool: Extent Pool P1 successfully created.
dscli> lsextpool
First half of the command output:
Name    ID stgtype rankgrp status availstor (2^30B) %allocated available reserved numvols
==========================================================================================
FB_0    P0 fb            0   full                 0        100         0        0       0
FB_1    P1 fb            1   full                 0        100         0        0       0

Second half of the command output:
Name    totalLogicalCap(MiB/cyl) usedLogicalCap(MiB/cyl) provisionedCap(MiB/cyl) totalPhysicalCap(MiB/cyl) usedPhysicalCap(MiB/cyl)
=============================================================================================================================
FB_0    0                        0                       0                       0                        0
FB_1    0                        0                       0                       0                        0

dscli> chrank -extpool P0 R0
CMUC00008I chrank: Rank R0 successfully modified.
dscli> chrank -extpool P1 R1
CMUC00008I chrank: Rank R1 successfully modified.
dscli> lsextpool
First half of the command output:
Name    ID stgtype rankgrp status availstor (2^30B) %allocated available reserved numvols
==========================================================================================
FB_0    P0 fb            0  below            101000          0   6463982       64       0
FB_1    P1 fb            1  below            101000          0   6463982       64       0

Second half of the command output:
Name    totalLogicalCap(MiB/cyl) usedLogicalCap(MiB/cyl) provisionedCap(MiB/cyl) totalPhysicalCap(MiB/cyl) usedPhysicalCap(MiB/cyl)
=============================================================================================================================
FB_0    108135060                1477436                 40960                   23605023                 1472246
FB_1    108135060                1477436                 40960                   23605023                 1472246
```

After a rank is assigned to an extent pool, use the **lsrank** command to display rank assignments to extent pools. In Example 10-32, you can see that rank R0 is assigned to extpool P0 and rank R1 to extent pool P1.

Example 10-32 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
```

To display more information about ranks, use **lsrank -l** as shown in Example 10-33. In addition to the extent size, the output includes more information on capacity and compression. In Example 10-33, **compressionType** has value *compression* because these ranks include FlashCore Modules disks and the **compressionRatio** value changes when the volumes are created and in use.

Example 10-33 List rank command showing the extent size, capacity, and compression

```
dscli> lsrank -l
First half of the command output:
ID Group State  datastate Array RAIDtype extpoolID extpoolnam stgtype exts    usedexts keygrp marray extsize
=============================================================================================================
R0      0 Normal Normal    A0           6 P0        FB_0       fb      6464054        0      1 MA1    16MiB
R1      1 Normal Normal    A1           6 P1        FB_1       fb      6464054        0      1 MA2    16MiB

Second half of the command output:
ID totalLogicalCap(MiB/cyl) usedLogicalCap(MiB/cyl) totalPhysicalCap(MiB/cyl) usedPhysicalCap(MiB/cyl) compressionRatio compressionType
=============================================================================================================================================
R0     108135060               1477304               23605023                1472246                1.00            compression
R1     108135060               1477304               23605023                1472252                1.00            compression
```

10.3.5 Creating the FB volumes

After ranks are associated to extents pools, 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.

An FB-type volume can be created as a fully provisioned volume, which is known as a standard volume or as a thin provisioned volume, which is referred as extent space efficient (ESE) volume.

Note: Only thin provisioned ESE volumes can be created in extent pools with FlashCore Modules (FCM) disks.

Extent pools with standard NVMe disks support both fully provisioned standard and thin provisioned ESE volumes.

Creating 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 a 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 DS8A00 can support 255 LSSs, and each LSS can support a maximum of 256 volumes.

Example 10-34 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 10-34 Create fixed-block volumes by using mkfbvol

```
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
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 10-34 on page 359, 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 10-34 on page 359, 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 2^{30} 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 10-34 on page 359 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 10-35, 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 10-35 Checking the machine after the volumes are created by using lsfbvol and lsextpool

```
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 10-36.

Default allocation policy: For DS8A00, the default allocation policy is rotate extents.

Example 10-36 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 10-37 on page 362) 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 previous examples, the extent pool P2 now consists of two ranks, R2 and R3.

Example 10-37 Using showfbvol for information about a striped volume

```
dscli> 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                        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                     PG0
migratingfrom               -
resgrp                      RG0
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       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 10-38.

Example 10-38 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 10-39. See both examples and check the difference.

Example 10-39 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               RG0
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.

Creating thin-provisioned ESE volumes

The DS8A00 supports thin-provisioned volumes. A thin-provisioned volume is created by the DS CLI **mkfbvol command** by specifying **ESE** storage allocation method (**SAM**) parameter as shown in Example 10-40.

Example 10-40 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 with 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 DS8000 Thin Provisioning (Updated for Release 9.3)*, 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.

Note: You cannot delete the volume even with the **force** parameter if the volume is in an active SafeGuarded copy relationship (GDPS or CSM). You must first remove the Safeguarded copy relationship before you can delete a volume.

Example 10-41 shows the creation of volumes 2200 and 2201, and then the assignment of volume 2200 to a volume group. An attempt to delete both volumes with the **-safe** option fails because the volume 2200 is assigned to volume group V0. 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 10-41 Deleting an FB volume

```
dscli> 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.
dscli> chvolgrp -action add -volume 2200 v0
CMUC00031I chvolgrp: Volume group V0 successfully modified.
dscli> rmfbvol -quiet -safe 2200-2201
CMUC00253E rmfbvol: Volume IBM.2107-75xxxxx/2200 is assigned to a user-defined
volume group. No volumes were deleted.
dscli> rmfbvol -quiet -safe 2201
CMUC00028I rmfbvol: FB volume 2201 successfully deleted.
dscli> 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 10-42):

- 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 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 for the command to proceed with the reinitialization process.

Example 10-42 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 10-43 shows the results of each command.

Example 10-43 Listing the host types by running the lshostype 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 10-44, the example host type is AIX with the address discovery method `SCSI Mask`.

Example 10-44 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 previously created volume group V1.

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 10-45 shows adding volume 2003 to volume group V1, displaying the results, and then removing the volume.

Example 10-45 Changing a volume group by using chvolgrp

```
dscli> chvolgrp -action add -volume 2003 V1
CMUC00031I chvolgrp: Volume group V1 successfully modified.
dscli> showvolgrp V1
Name AIX_VG_01
ID   V1
Type SCSI Mask
Vols 2000 2001 2002 2003 2100 2101 2102
dscli> chvolgrp -action remove -volume 2003 V1
CMUC00031I chvolgrp: Volume group V1 successfully modified.
dscli> 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 280.

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 with 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 10-46.

Example 10-46 Creating generic host types Linux Server and Windows Server

```
dscli> 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 -
x355O_2_03 Windows Server  Online -
x365O_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 10-47 provides examples of the commands.

Example 10-47 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 using the **mkhostport** command. This task can also be done with the **mkhost** command specifying 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 displays 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 10-48 shows the host port connection.

Example 10-48 Host port connection

```
dscli> lshostport -login
WWNN             WWPN             Host       I/O ports ESSIOport
============================================================
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
dscli> lshostport -unknown
Date/Time: October 16, 2025 at 4:47:19 PM EDT IBM DSCLI Version: 7.10.10.348 DS:
IBM.2107-75xxxxx
WWNN             WWPN             ESSIOport
=========================================
20000000C9F607E8 10000000C9F607E8 I0032
```

```
20000000C9F607E9 10000000C9F607E9 I0103
20000120FA13B914 10000090FA13B914 I0032
20000120FA13B915 10000090FA13B915 I0103
dscli> mkhostport -host Host_1 "10000000C9F607E8"
CMUC00542I mkhostport: The host port 10000000C9F607E8 is successfully created.
dscli> mkhostport -host Host_1 "10000000C9F607E9"
CMUC00542I mkhostport: The host port 10000000C9F607E9 is successfully created.
dscli> 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
dscli> 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 10-49 shows the steps for a cluster and a host.

Example 10-49 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 10-50 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 10-50 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.2107-75xxxxx/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 10-44 on page 367 shows the already created volume group for an example AIX host connection. The volumes are already assigned to the volume group.

Example 10-51 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 10-43 on page 366. 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 10-51 Creating host connections by using mkhostconnect and lshostconnect

```
dscli> lshostport -unknown
WWNN             WWPN             ESSIOport
=========================================
20000120FA13B914 10000090FA13B914 I0032
20000120FA13B915 10000090FA13B915 I0103
dscli> 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
dscli> showvolgrp V1
Name AIX_VG_01
ID   V1
Type SCSI Mask
Vols 2000 2001 2002 2100 2101 2102
dscli> mkhostconnect -wwname 10000090FA13B914 -hosttype pSeries -volgrp V1 AIX_Server_01
CMUC00012I mkhostconnect: Host connection 000A successfully created.
dscli> mkhostconnect -wwname 10000090FA13B915 -hosttype pSeries -volgrp V1 AIX_Server_01
CMUC00012I mkhostconnect: Host connection 000B successfully created.
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
```

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 DS8A00, 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.

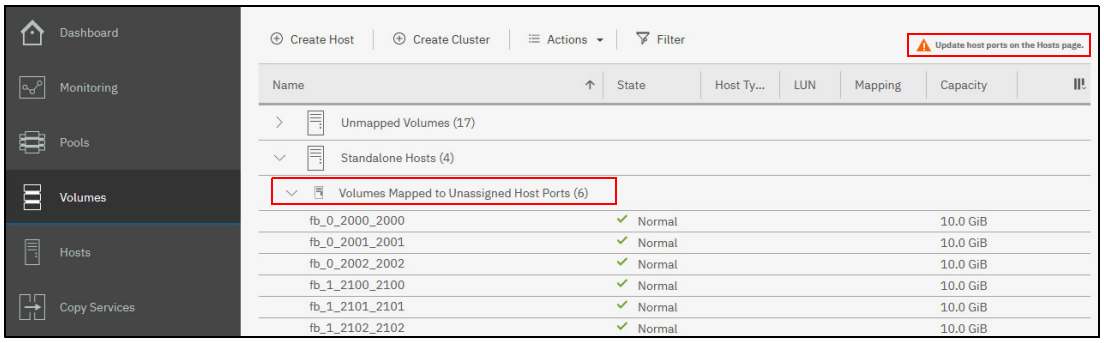


Figure 10-8 Volumes that are mapped to a host port without a host

The **lshostconnect -l** command in Example 10-52 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 10-52 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 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 10-53 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 10-53 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 10-54 with the command **lsconnect -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 10-54 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 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 DS8000 storage configuration for the CKD volumes

This list contains the steps to configure CKD storage in the DS8000:

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.

Unlike with FB 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 Disk classes

More information about available disk classes in DS8A00 can be found in 10.3.1, “Disk classes” on page 354.

10.4.2 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 354.

10.4.3 Creating the ranks and extent pools

After you create all of the required arrays, create the ranks by using the **mkrank** command. The format of the command for creating Count Key Data (CKD) rank is **mkrank -array A*x* -stgtype CKD**.

If you enabled Data-at-Rest encryption with local or external key manager, then when you create ranks, you must include the **-keygrp** flag in the **mkrank** command with the corresponding *Encryption Key Group* **ID** parameter which is included in the **lskeygrp** command output.

After all the ranks are created, use the **lsrank** command to display the following information:

- List of all created ranks with associated ID.
- The rank group 0 or 1. Each group has server affinity (0 for server0 or 1 for server1). Rank group ID is displayed only when a rank is assigned to a specific extent pool.
- The RAID type.

The format of the rank (fb or ckd).

Example 10-55 shows the process of creating ranks using **lskeygrp**, **mkrank**, **lsrank**, **mkextpool**, **chrank**, and **lsextpool** commands.

Example 10-55 Steps to create CKD ranks

```
dscli> lskeygrp
ID  state      reckeystate reckeydate datakeydate keyprotocol type name
==========================================================================
1   accessible disabled    -          09/18/2024  LOCAL       DAR  DAR_1

As per the above output, the Encryption Key Group ID = 1. It must be specified in
the mkrank command in -keygrp parameter.

dscli> mkrank -array A0 -keygrp 1 -stgtype ckd
CMUC00007I mkrank: Rank R0 successfully created.
dscli> mkrank -array A1 -keygrp 1 -stgtype ckd
CMUC00007I mkrank: Rank R1 successfully created.

dscli> lsrank
ID Group State      datastate Array RAIDtype extpoolID stgtype
============================================================
R0     - Unassigned Normal     A0          6 -         ckd
R1     - Unassigned Normal     A1          6 -         ckd

dscli> lsarray
Array State    Data   RAIDtype    arsite Rank DA Pair DDMcap (10^9B)
```

```
==============================================================
A0    Assigned Normal 6 (5+P+Q+S)   S1     R0  0              4800.0
A1    Assigned Normal 6 (5+P+Q+S)   S2     R1  0              4800.0
```

When you define 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 entire extent pool. Example 10-56 shows CKD ranks with small and large extent sizes.

Example 10-56 CKD ranks with different extent sizes

```
dscli> lsrank -l
ID Group State  datastate Array RAIDtype extpoolID extpoolnam stgtype exts   usedexts keygrp marray extsize (cap)
=========================================================================================================================
R10    0 Normal Normal    A10          6 P2        CKD_FCM_0  ckd     8099          0      1 MA1    1113cyl
R11    1 Normal Normal    A11          6 P3        CKD_FCM_1  ckd     429258        0      1 MA2    21cyl
```

10.4.4 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 specified by the **-stgtype** parameter and for IBM z mainframe is **CKD**.
- 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 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 a best practice to note the ID.

The **mkextpool** command forces you to name the extent pools. To do so, complete the following steps (see Example 10-57 on page 378):

1. Create empty extent pools by using the **mkextpool** command. The **keygrp** parameter must be used because encryption is enabled (see Example 10-30 on page 356 for more information in finding the encryption group ID value).
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 **lsextpool** and note the change in the capacity of the extent pool after assigning ranks into the extent pools.

Example 10-57 Create CKD Extent Pools and assign ranks to Extent Pools

```
dscli> mkextpool -rankgrp 0 -keygrp 1 -stgtype ckd CKD_0
CMUC00000I mkextpool: Extent pool P0 successfully created.
dscli> mkextpool -rankgrp 1 -keygrp 1 -stgtype ckd CKD_1
CMUC00000I mkextpool: Extent Pool P1 successfully created.
dscli> lsextpool
First half of the command output:
Name    ID stgtype rankgrp status availstor (2^30B) %allocated available reserved numvols
==========================================================================================
CKD_0   P0 ckd           0   full                 0         100         0        0       0
CKD_1   P1 ckd           1   full                 0         100         0        0       0

Second half of the command output:
Name    totalLogicalCap(MiB/cyl) usedLogicalCap(MiB/cyl) provisionedCap(MiB/cyl) totalPhysicalCap(MiB/cyl) usedPhysicalCap(MiB/cyl)
=============================================================================================================================
CKD_0   0                        0                       0                       0                        0
CKD_1   0                        0                       0                       0                        0

dscli> chrank -extpool P0 R0
CMUC00008I chrank: Rank R0 successfully modified.
dscli> chrank -extpool P1 R1
CMUC00008I chrank: Rank R1 successfully modified.
dscli> lsextpool
First half of the command output:
Name    ID stgtype rankgrp status availstor (2^30B) %allocated available reserved numvols
==========================================================================================
CKD_0   P0 ckd           0   below            101000          0   6463982       64       0
CKD_1   P1 ckd           1   below            101000          0   6463982       64       0

Second half of the command output:
Name    totalLogicalCap(MiB/cyl) usedLogicalCap(MiB/cyl) provisionedCap(MiB/cyl) totalPhysicalCap(MiB/cyl) usedPhysicalCap(MiB/cyl)
=============================================================================================================================
CKD_0   108135060                1477436                 40960                   23605023                 1472246
CKD_1   108135060                1477436                 40960                   23605023                 1472246
```

After a rank is assigned to an extent pool, use **lsrank** command to display rank assignments to extent pools. In Example 10-58, you can see that rank R0 is assigned to extpool P0 and rank R1 to extent pool P1.

Example 10-58 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        ckd
R1      1 Normal Normal    A1           6 P1        ckd
```

In order to display more information about ranks, use **lsrank -**l command as shown in Example 10-59. In addition to the extent size, the output includes more information on capacity and compression. In Example 10-59, **compressionType** has value *compression* because these ranks include FlashCore Modules disks and **compressionRatio** value changes when the volumes are created and in use.

Example 10-59 List ranks with **lsrank -l** *command*

```
dscli> lsrank -l
First half of the comand output:
ID Group State  datastate Array RAIDtype extpoolID extpoolnam stgtype exts    usedexts keygrp marray extsize
=============================================================================================================
```

```
R0     0 Normal Normal    A0          6 P0       CKD_0      ckd     6464054       0      1 MA1    21cyl
R1     1 Normal Normal    A1          6 P1       CKD_1      ckd     6464054       0      1 MA2    21cyl
```

Second half of the command output:

```
ID totalLogicalCap(MiB/cyl) usedLogicalCap(MiB/cyl) totalPhysicalCap(MiB/cyl) usedPhysicalCap(MiB/cyl) compressionRatio compressionType
==========================================================================================================================================
R0     108135060               1477304               23605023                1472246                1.00            compression
R1     108135060               1477304               23605023                1472252                1.00            compression
```

10.4.5 Logical control unit creation

In a CKD environment, you must create LCUs before the volumes are created. In Example 10-60, you can see what happens if you try to create a CKD volume without creating an LCU first.

Example 10-60 Trying to create CKD volumes without creating an LCU first

```
dscli> 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 **mklcu** command. The command uses the following format:

```
mklcu -qty XX -id XX -ss XXXX
```

Note: For the z/OS hardware definition and copy services (Metro Mirror or Global Mirror), the subsystem identifier (SSID) (-ss -id) must be unique for all connected storage systems.

To display the LCUs that you created, run the **lslcu** command.

Example 10-61 shows the creation of two LCUs by running the **mklcu** command and then listing the created LCUs by running the **lslcu** command. By default, the LCUs that were created are the 3990-6 type.

Example 10-61 Creating a logical control unit by running mklcu

```
dscli> mklcu -qty 2 -id BC -ss 91BC
CMUC00017I mklcu: LCU BC successfully created.
CMUC00017I mklcu: LCU BD successfully created.
dscli> lslcu
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 can be maximized by spreading the workload across both servers in the DS8A00.

Important: For the DS8A00, 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 DS8A00. Each LCU can manage a maximum of 256 volumes, including alias volumes for the parallel access volume (PAV) feature.

10.4.6 Creating the CKD volumes

When LCUs are created, you can start creating CKD volumes. A CKD-type volume can be created as a fully provisioned volume, which is known as a standard volume or as a thin provisioned volume, which is referred to as an extent space efficient (ESE) volume.

Note: Only thin provisioned ESE volumes can be created in extent pools with FlashCore Modules (FCM) disks.

Extent pools with standard NVMe disks support both fully provisioned standard and thin provisioned ESE volumes.

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* `-name` *CKD_EAV1_#h BC06*

With CKD volumes, the capacity is expressed in cylinders or as mod1 (Model 1) extents (1113 cylinders). To use the capacity more efficiently and avoid wasting space, use volume capacities that are a multiple of 1113 cylinders.

The DS8A00 supports EAV volumes up to 1,182,006 cylinders. The EAV device type is called *3390 Model A* and must be specified in the **mkckdvol** command using parameter **datatype 3390-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 match the control unit address (CUADD) 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 10-62 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 10-62 Creating CKD volumes by using mkckdvol

```
dscli> mkckdvol -extpool P2 -cap 262668 -datatype 3390-A -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.
dscli> 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 DS8A00 supports thin-provisioned volumes. A thin-provisioned volume is created with the **mkckdvol** using the `sam ese` parameter, as shown in Example 10-63.

Example 10-63 Creating a thin-provisioned CKD volume

```
dscli> 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, SPACERel, with its options. A corresponding IBM RACF FACILITY class profile provides protection as well.

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 DS8000 Thin Provisioning (Updated for Release 9.3)*, 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*.

Note: Rotate extents is the default allocation policy.

The `eam` (extent allocation method) parameter is specified with `rotateexts` or `rotatevols` option in the **`mkckdvol`** command (Example 10-64).

Example 10-64 Creating a CKD volume with extent pool striping

```
dscli> mkckdvol -extpool p3 -cap 10017 -name CKD_EAV1_#h -sam ese -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 10-65) 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 10-65 Obtaining information about a striped 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                 477
cap (cyl)            10017
cap (10^9B)          8.5
cap (2^30B)          7.9
ranks                2
sam                  ese
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
safeguardedcapstate     -
totalLogicalCap         151298364
usedLogicalCap          0
usedPhysicalCap         0
compressionType         compression
==============Rank extents==============
rank extents capacity (MiB/cyl) metadata
========================================
R2       239 5019               yes
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 10-66. The new capacity must be larger than the previous capacity. *You cannot shrink the volume*.

Example 10-66 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 10-67. In this example, the pool `P3` is using small 21-cylinder CKD extents.

Example 10-67 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                      ese
repcapalloc              -
eam                      rotateexts
reqcap (cyl)             30051
cap (Mod1)               27.0
realextents              1431
virtualextents           4
realcap (cyl)            30051
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
safeguardedcapstate      -
totalLogicalCap          151298364
usedLogicalCap           0
usedPhysicalCap          0
compressionType          compression
==============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 Copy Services 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.

Example 10-68 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                     ese
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
safeguardedcapstate     -
totalLogicalCap         151298364
usedLogicalCap          0
usedPhysicalCap         0
compressionType         compression
==============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                     ese
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
safeguardedcapstate     -
totalLogicalCap         151298364
usedLogicalCap          0
usedPhysicalCap         0
compressionType         compression
==============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.

The command includes the capability to prevent the accidental deletion of volumes that are in use. A CKD volume is considered *in use* if it participates in a copy services 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 in the command, the volumes are deleted without checking to see whether they are in use.

Note: You cannot delete the volume even with the `force` parameter if the volume is in an active SafeGuarded copy relationship (GDPS or CSM). You must first terminate Safeguarded copy relationship before you can delete a volume with `force`.

Example 10-69 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 copy services relationship. The **`rmckdvol BD02-BD03`** command deletes only volume BD03, which is offline. To delete volume BD02, use the `-force` parameter.

Example 10-69 Deleting CKD volumes

```
dscli> 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

dscli> 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.

dscli> 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

dscli> 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.

dscli> 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 Copy Services 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 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 for the command to proceed with the reinitialization process, as shown in Example 10-70.

Example 10-70 An initckdvol command example

```
dscli> 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.
```

10.4.7 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 copy services functions only.

For more information about resource groups, see *IBM System Storage DS8000 Copy Services Scope Management and Resource Groups*, REDP-4758.

10.4.8 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 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 publication:

- *IBM DS8000 Easy Tier (Updated for DS8000 R9.0)*, REDP-5667

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 9.15, “Performance monitoring” on page 319.

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.

10.5.1 Overview of metrics

Example 10-71 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 10-71 Metrics for a specific fixed-block volume

```
dscli> showfbvol -metrics 000A
ID                       000A
Date                     10/10/2024 14:19:03 CEST
normrdrqts               7129645
normrdhits               4725660
normwritereq             25533232
normwritehits            25533232
seqreadreqs              11746083
seqreadhits              11578653
seqwritereq              3994518
seqwritehits             3994518
cachfwrreqs              0
cachfwrhits              0
cachefwreqs              0
cachfwhits               0
inbcachload              0
bypasscach               0
DASDtrans                4019246
seqDASDtrans             5825774
cachetrans               14933989
NVSspadel                0
normwriteops             0
seqwriteops              0
reccachemis              1314738
qwriteprots              0
CKDirtrkac               0
CKDirtrkhits             0
cachspdelay              0
timelowifact             0
thread                   9848943
phwrite                  820103
phbyteread               4269617
phbytewrite              7453693
recmoreads               2104182
sfiletrkreads            0
contamwrts               0
PPRCtrks                 0
NVSspallo                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 10-72 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 10-72 Metrics for a specific CKD volume

```
dscli> showckdvol -metrics 7b3d
ID                         7B3D
normrdrqts                 9
normrdhits                 9
normwritereq               0
normwritehits              0
seqreadreqs                0
seqreadhits                0
seqwritereq                0
seqwritehits               0
cachfwrreqs                0
cachfwrhits                0
cachefwreqs                0
cachfwhits                 0
inbcachload                0
bypasscach                 0
DASDtrans                  201
seqDASDtrans               0
cachetrans                 1
NVSspadel                  0
normwriteops               0
seqwriteops                0
reccachemis                0
qwriteprots                0
CKDirtrkac                 9
CKDirtrkhits               9
cachspdelay                0
timelowifact               0
phread                     201
phwrite                    1
phbyteread                 49
phbytewrite                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 10-73 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 10-73 Metrics for a specific rank

```
dscli> showrank -metrics R0
ID            R0
Date          10/10/2024 14:32:46 CEST
byteread      8175947
bytewrit      29508105
Reads         17783022
Writes        17946281
timeread      1105024
timewrite     19434052
dataencrypted yes
timeload      8206273
dscli>
```

Example 10-74 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 DS8A00 exchange information. If one of these components has problems because of hardware or configuration issues, all of the other components are affected.

Example 10-74 Metrics for a specific 32 Gbps I/O port

```
dscli> showioport -metrics I0032
ID                       I0032
Date                     10/12/2024 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 10-75 on page 394 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 (In-flight encryption over FICON links).

The following groups of errors point to specific problem areas:

- Any nonzero figure in the counters `LinkFailErr`, `LossSyncErr`, `LossSigErr`, and `PrimSeqErr` indicates that unstable HBAs 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 DS8A00. 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 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 are used 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.
    - DS8A00 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 10-75 provides an example of the RDP status.

Example 10-75 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             32GFC
Advertised B-B Credit            20
Attached Port B-B Credit         180
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               10/12/2024 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             32GFC
Advertised B-B Credit            20
Attached Port B-B Credit         180
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               10/12/2024 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>
```

10.5.2 Offloading 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:

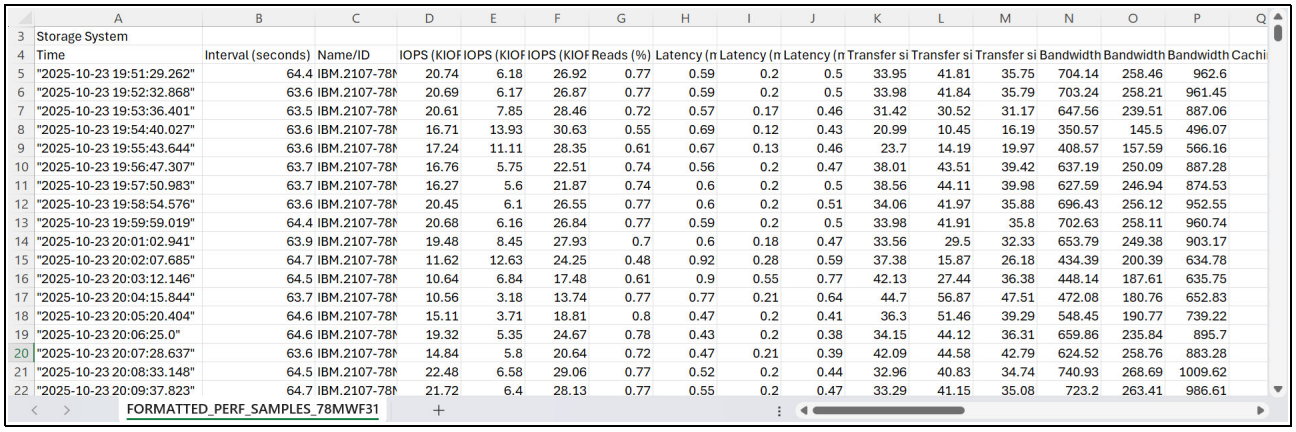
-auditlog	Exports one file that contains the audit log for the Management Console (MC) server.
-sdocert	Exports the IBM Certified Secure Data Overwrite (SDO) certificate for each storage facility.
-etdata	Exports two files that contain the Easy Tier summary data.
-config	Offloads two files: One file 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 10-76 is a *.csv file* with detailed information. For more information, see Figure 10-9 on page 396.

Example 10-76 The offloadfile command that is issued from DS CLI

```
dscli> offloadfile -perfdata c:\temp
CMUC00428I offloadfile: The perfdata file has been offloaded to
c:\temp\FORMATTED_PERF_SAMPLES_75xxxxx.csv.
dscli>
```



	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
3	Storage System																
4	Time	Interval (seconds)	Name/ID	IOPS (KIOF	IOPS (KIOF	IOPS (KIOF	Reads (%)	Latency (n	Latency (n	Latency (n	Transfer si	Transfer si	Transfer si	Bandwidth	Bandwidth	Bandwidth	Cachi
5	"2025-10-23 19:51:29.262"	64.4	IBM.2107-78N	20.74	6.18	26.92	0.77	0.59	0.2	0.5	33.95	41.81	35.75	704.14	258.46	962.6	
6	"2025-10-23 19:52:32.868"	63.6	IBM.2107-78N	20.69	6.17	26.87	0.77	0.59	0.2	0.5	33.98	41.84	35.79	703.24	258.21	961.45	
7	"2025-10-23 19:53:36.401"	63.5	IBM.2107-78N	20.61	7.85	28.46	0.72	0.57	0.17	0.46	31.42	30.52	31.17	647.56	239.51	887.06	
8	"2025-10-23 19:54:40.027"	63.6	IBM.2107-78N	16.71	13.93	30.63	0.55	0.69	0.12	0.43	20.99	10.45	16.19	350.57	145.5	496.07	
9	"2025-10-23 19:55:43.644"	63.6	IBM.2107-78N	17.24	11.11	28.35	0.61	0.67	0.13	0.46	23.7	14.19	19.97	408.57	157.59	566.16	
10	"2025-10-23 19:56:47.307"	63.7	IBM.2107-78N	16.76	5.75	22.51	0.74	0.56	0.2	0.47	38.01	43.51	39.42	637.19	250.09	887.28	
11	"2025-10-23 19:57:50.983"	63.7	IBM.2107-78N	16.27	5.6	21.87	0.74	0.6	0.2	0.5	38.56	44.11	39.98	627.59	246.94	874.53	
12	"2025-10-23 19:58:54.576"	63.6	IBM.2107-78N	20.45	6.1	26.55	0.77	0.6	0.2	0.51	34.06	41.97	35.88	696.43	256.12	952.55	
13	"2025-10-23 19:59:59.019"	64.4	IBM.2107-78N	20.68	6.16	26.84	0.77	0.59	0.2	0.5	33.98	41.91	35.8	702.63	258.11	960.74	
14	"2025-10-23 20:01:02.941"	63.9	IBM.2107-78N	19.48	8.45	27.93	0.7	0.6	0.18	0.47	33.56	29.5	32.33	653.79	249.38	903.17	
15	"2025-10-23 20:02:07.685"	64.7	IBM.2107-78N	11.62	12.63	24.25	0.48	0.92	0.28	0.59	37.38	15.87	26.18	434.39	200.39	634.78	
16	"2025-10-23 20:03:12.146"	64.5	IBM.2107-78N	10.64	6.84	17.48	0.61	0.9	0.55	0.77	42.13	27.44	36.38	448.14	187.61	635.75	
17	"2025-10-23 20:04:15.844"	63.7	IBM.2107-78N	10.56	3.18	13.74	0.77	0.77	0.21	0.64	44.7	56.87	47.51	472.08	180.76	652.83	
18	"2025-10-23 20:05:20.404"	64.6	IBM.2107-78N	15.11	3.71	18.81	0.8	0.47	0.2	0.41	36.3	51.46	39.29	548.45	190.77	739.22	
19	"2025-10-23 20:06:25.0"	64.6	IBM.2107-78N	19.32	5.35	24.67	0.78	0.43	0.2	0.38	34.15	44.12	36.31	659.86	235.84	895.7	
20	"2025-10-23 20:07:28.637"	63.6	IBM.2107-78N	14.84	5.8	20.64	0.72	0.47	0.21	0.39	42.09	44.58	42.79	624.52	258.76	883.28	
21	"2025-10-23 20:08:33.148"	64.5	IBM.2107-78N	22.48	6.58	29.06	0.77	0.52	0.2	0.44	32.96	40.83	34.74	740.93	268.69	1009.62	
22	"2025-10-23 20:09:37.823"	64.7	IBM.2107-78N	21.72	6.4	28.13	0.77	0.55	0.2	0.47	33.29	41.15	35.08	723.2	263.41	986.61	

FORMATTED_PERF_SAMPLES_78MWF31

Figure 10-9 Snapshot from the -perfdata csv file

10.6 Private network security commands

There are DS CLI commands used to manage DS8A00 network security. 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 DS8A00. 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 HMC settings:

  - Enable and disable the HMC command-line shell access
  - Enable and disable the HMC WUI access
  - Enable and disable Assist On-site (AOS) or IBM Remote Support Center (RSC) access

  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 10-77.

Example 10-77 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 10-78 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 10-78 Enabling TLS for secure syslog traffic

```
dscli> mksyslogserver -addr 9.100.10.7 -protocol tls -cacert
/Users/heping/ut/syslog/ca.pem -hmccert /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 10-79.

Example 10-79 The lssyslogserver -l command and a sample output

```
dscli> 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 external encryption key servers (typically IBM Guardium 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 10-80.

Example 10-80 The managekeygrp command for importing a custom certificate

```
dscli> 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 10-81. For more information about the encryption and DS8000 certificates, see *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security*, REDP-4500.

Example 10-81 The managekeygrp command for updating a custom certificate

```
dscli> managekeygrp -action updatecert -certType CUSTOMER -key data 1
CMUC00472I managekeygrp: The certificate for encryption group 1 has been updated.
```

For more information, see IBM Documentation.

10.7 Copy Services commands

More DS CLI commands are available, and many of these commands are used for the management of copy services, such as the Safeguarded copy, FlashCopy, Metro Mirror (MM), and Global Mirror (GM) commands.

These commands are not described in this chapter. For more information, see Copy Services commands.

10.8 Earlier DS CLI commands and scripts

As versions of the DS8000 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 DSI CLI **help** command, as shown in Example 10-82.

Example 10-82 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 10-83.

Example 10-83 Support for older commands

```
dscli> lsddm
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.	**lsstgencl**
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
lspciecard	Lists PCIe cards.	N/A

10.9 Dual Control commands

IBM Storage DS8000® (10th Generation), Version 10 Release 1 (10.1.0), introduces a new Dual Control feature across the DS8000 GUI, DSCLI, and REST API interfaces. Dual Control is a security enhancement that requires two separate authorized individuals to approve and execute critical operations. This follows the Maker/Checker principle: a Maker submits a request for a protected operation, and a Checker must approve the request before execution.

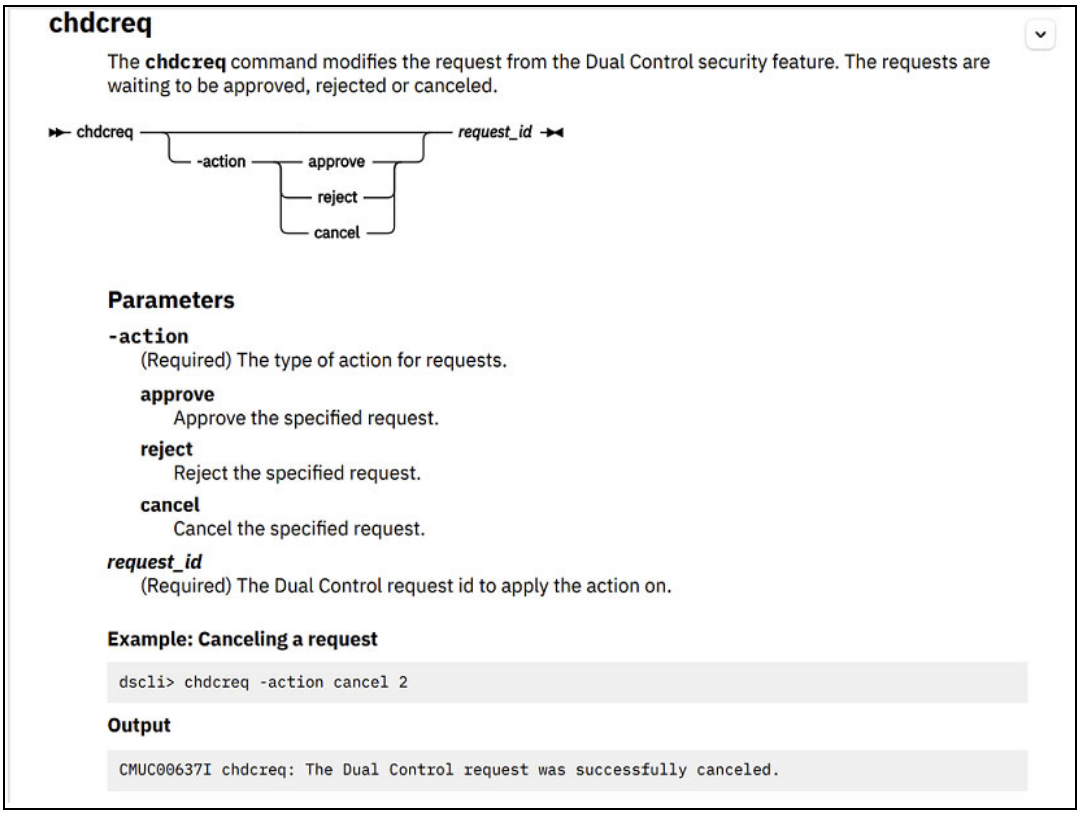
Dual Control includes a setting to enable or disable Dual Control Mode. When enabled, a designated Checker must authorize the task. This feature helps prevent data loss caused by intentional or accidental errors, protects against privilege escalation, and ensures that critical tasks require multiple authorized approvals - while maintaining productivity for non-critical operations.

10.9.1 Examples of Dual Control DSCLI commands

The following are examples of Dual Control DSCLI commands.

chdcreq

The **chdcreq** command modifies the request from the Dual Control security feature. The requests are waiting to be approved, rejected or canceled. See Figure 10-10 on page 400.



chdcreq

The **chdcreq** command modifies the request from the Dual Control security feature. The requests are waiting to be approved, rejected or canceled.

Parameters

-action
(Required) The type of action for requests.

approve
Approve the specified request.

reject
Reject the specified request.

cancel
Cancel the specified request.

request_id
(Required) The Dual Control request id to apply the action on.

Example: Canceling a request

```
dscli> chdcreq -action cancel 2
```

Output

```
CMUC00637I chdcreq: The Dual Control request was successfully canceled.
```

Figure 10-10 chdcreq command

chuser

The **chuser** command modifies and locks or unlocks a DS CLI or a DS8000 Storage Management GUI basic user account. See Figure 10-11 on page 401.

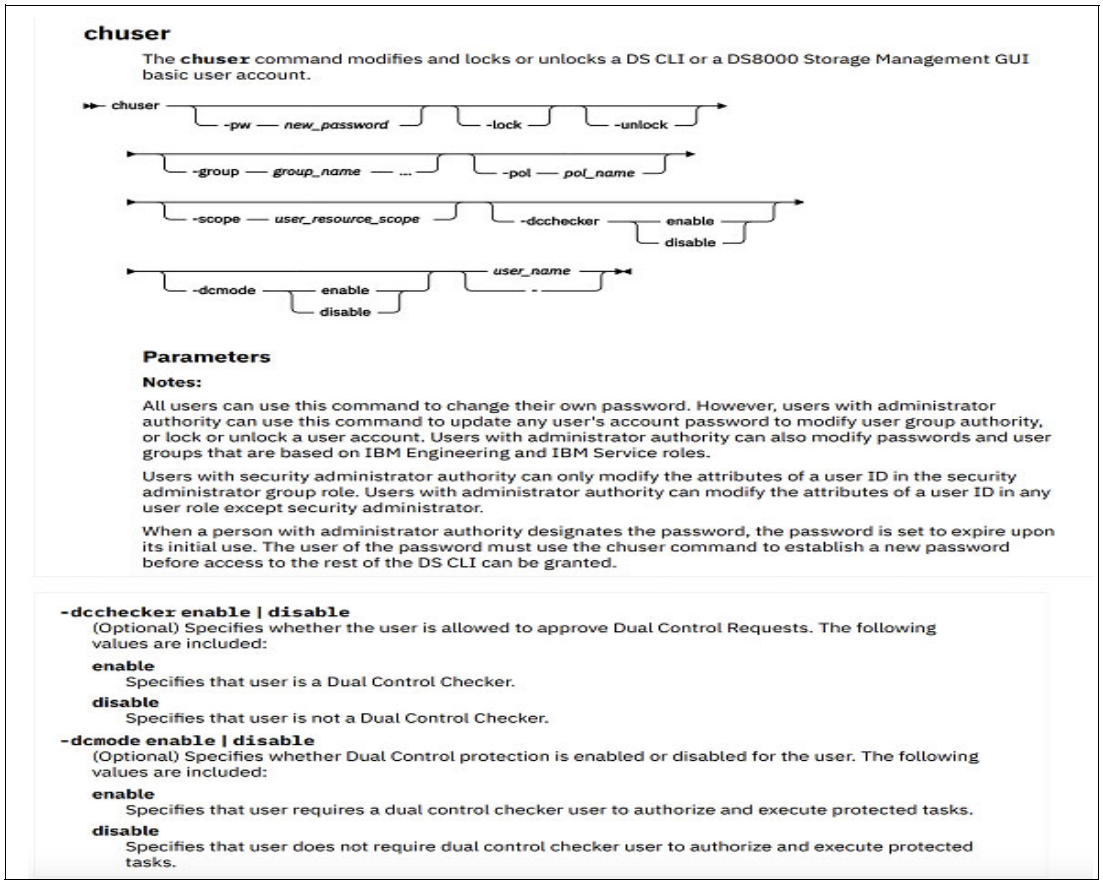


Figure 10-11 chuser command

lsdcreq

The **lsdcreq** command displays a list of request from the Dual Control security feature and status information for each request in the list. See Example 10-84 and Figure 10-12 on page 402.

Example 10-84 Listing Dual Control requests

```
dscli> lsdcreq -l

ID Request Date

================================================

1 Create Administrator 2025-07-08 18:59:16.625

2 Create Administrator 2025-07-08 19:04:44.508

User Description

================================================

admin user testOAdmin with role Administrator

dc_maker user testAdmin with role Administrator
```

```
dscli> lsdcreq
Date/Time: August 4, 2025 at 2:05:38 PM MST IBM DSCLI Version: 7.10.10.281 DS: -
ID Request                                    State     Date                    User
================================================================================
1  Delete Volumes                             Submitted 2025-08-04 13:37:46.045 admin
2  Create Administrator                       Submitted 2025-08-04 13:38:11.779 admin
3  Remove DC protection from user             Submitted 2025-08-04 13:38:45.384 admin
4  Allow Checker Authority                    Submitted 2025-08-04 13:40:46.143 admin
5  Modify Remote Support Center Connection Submitted 2025-08-04 13:42:05.424 admin
dscli>
```

Figure 10-12 lsdcreq command

showdcreq

The **showdcreq** command displays detailed properties of the request from Dual Control security. See Figure 10-13.

showdcreq

The **showdcreq** command displays detailed properties of the request from Dual Control security.

Syntax



Parameters

***request_id* | -**
(Required) The ID of the Dual Control request. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Examples

Displaying details for user accounts

```
dscli> showdcreq 2
```

Output:

Note: For this command and all other DS CLI show commands, the results are shown in table format for clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output reports that are associated with the **showdcreq** command.

Date/Time: July 8, 2025 7:23:48 PM CDT IBM DSCLI Version: 0.0.0.0 DS: -

Field	Value
ID	1
Request	Enable Dual Control
Date	2025-07-08 19:23:45.813
User	dc_maker
Description	Enable Dual Control

Figure 10-13 showdcreq command

For more information related to Dual Control DSCLI commands see IBM DS8000 CLI Documentation.

10.10 For more information

For more information about the topics that were described in this chapter, see the following resources:

- For more information about the copy services configuration, see IBM DS8000 Series Command-Line Interface User's Guide.
- 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*, REDP-4500.
- For more information about thin-provisioning, see *IBM DS8000 Thin Provisioning*, REDP-5343.
- For more information about DS CLI commands that relate to Easy Tier, see *IBM DS8000 Easy Tier,* REDP-4667.

11

Licensed Machine Code

This chapter describes considerations that relate to the planning and installation of new Licensed Machine Code (LMC) bundles on the IBM DS8A00. The overall process for the DS8A00 is the same as for previous models.

This chapter covers the following topics:

- “How new Licensed Internal Code is released” on page 406
- “Bundle installation” on page 408
- “Code updates” on page 418
- “Host adapter firmware updates” on page 418
- “Loading the code bundle” on page 420
- “Postinstallation activities” on page 421
- “Summary” on page 421

© Copyright IBM Corp. 2025. All rights reserved.

11.1 How new Licensed Internal Code is released

Most of the hardware components within the DS8A00 system can be updated with new firmware when it is available:

- Device adapters (DAs)
- Host adapters (HAs)
- Flash Drives:
  - NVMe Industry Standard Flash Drives
  - NVMe FlashCore Modules (FCMs)
- Power subsystems:
  - Intelligent Power Distribution Units (iPDUs)
  - Storage Enclosure (SE) components:
    - Services Managers (ESMs)
    - Power Supplies (PSUs)
  - I/O Bay components:
    - I/O Bay logic cards
    - I/O Bay PSUs

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 DS8A00, new features are released through new LIC levels.

When IBM releases a new LIC for the DS8A00, it is released in the form of a bundle. The term *bundle* is used because a new code release can include updates for various DS8A00 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.

The cross-reference table at DS8A00 Code Bundle Information 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 current, maintain a current version of the Data Storage Command-line Interface (DS CLI).

The DS8A00 continues to use the same naming convention for bundles like the former DS8000 generations: `PR.MM.FFF.EEEE`.

The following list describes the components of the naming convention:

- P: Product (1 = DS8000)
- R: Release Major (*X*, 0= DS8A00)
- 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 on page 407.

Example 11-1 Bundle level information

```
For BUNDLE 10.10.106.0 :
Product          DS8000
Release Major    0
Release Minor    0
Fix Level        106
EFIX level       0
```

The 10.10 in Example 11-1 stands for the Release 10.1 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: October 24, 2025 at 8:44:42 AM CEST IBM DSCLI Version: 7.10.10.348 DS: -
DSCLI          7.10.10.348
StorageManager 10.2.2507181727
HMC DSCLI      7.10.10.348
================Version=================
Storage Image    LMC         Bundle Version
========================================
IBM.2107-78MWF31 7.10.10.348 10.10.106.0
```

The Bundle version (Release) can also be retrieved from the DS Storage Manager by clicking **Actions** → **Properties** from the Dashboard window, as shown in Figure 11-1 on page 408.

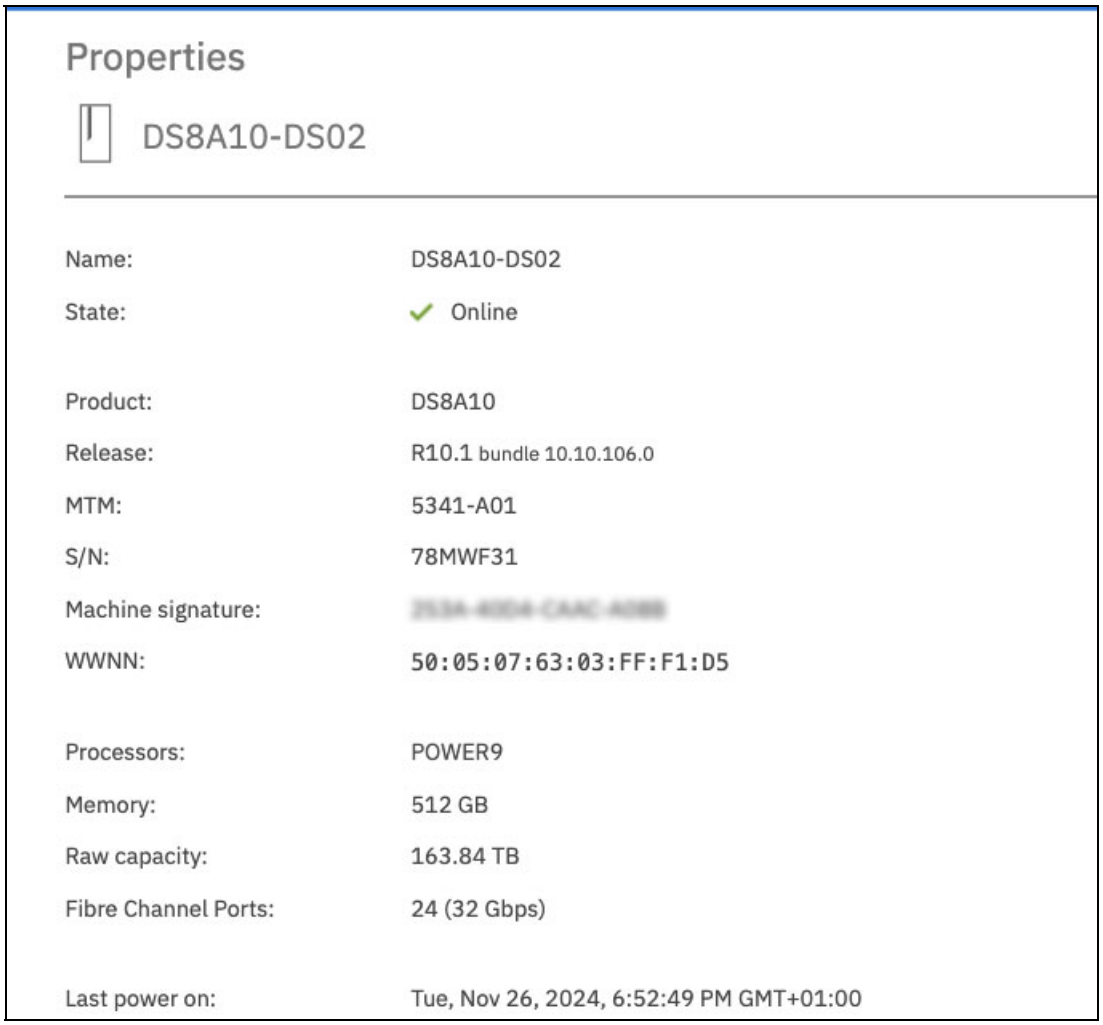


Figure 11-1 Bundle version under DS Storage Manager

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). The customer can also manage the entire process from the DS8000 Storage Manager GUI. During the planning phase customers must read the following document to know whether any prerequisites must be considered:

Customer "Must-Read" Information for DS8000 Code Updates

The bundle package contains the following 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
- HA code levels
- DA code level
- I/O bay code level
- Power code levels
- Flash Drive Enclosure Service Manager (ESM) interface card code levels
- Flash Drive enclosure power supply unit (PSU) code levels
- Flash Drive firmware code level

11.2.1 Code update methods

DS8000 LMC can be updated either remotely or locally at the HMC by an IBM Systems Service Representative (SSR):

- Remote Code Load (RCL): This is the standard method. IBM can perform RCL remotely, or the client can initiate the update through the DS8000 Storage Manager GUI.
- On-site Code Load (OCL): If the client prefers an IBM SSR to perform the update onsite, the appropriate Feature Code must be ordered with the system:
  - FC AHY2 for systems with Expert Care Premium
  - FC AHY3 for systems with Expert Care Advanced

Options by Expert Care level:

- Expert Care Premium (the default): Includes RCL (2x per year) by default. Customer-controlled code load remains available. Optional upgrade to OCL is possible with FC AHY2. Premium also includes IBM Storage Technical Account Manager (TAM) and Predictive Support with IBM Storage Insights Pro.
- Expert Care Advanced: Default method is customer-controlled code load. Clients can add:
  - RCL (2x per year) with FC AHY4
  - OCL by adding FC AHY3
- No Expert Care: Rare cases where DS8000 systems are configured without an Expert Care SLA. The only option is customer-controlled code load.

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 from IBM Fix Central on the HMC.

After the CDA software preinstallation starts, the following steps occur automatically:

1. The release bundle is downloaded from either the physical media or network to the Primary HMC (HMC1) flash disk drive (SSD).
2. The release bundle is copied from HMC1 to the Secondary HMC (HMC2)
3. The HMCs receive any code update-specific fixes.
4. Code updates are distributed to the logical partition (LPAR) and staged on an alternative base operating system (BOS) repository.

5. 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 HMC1.
2. Updates code on 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 DS8A00 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 4.3.2, “Host connections” on page 120.
5. New firmware for the hypervisor, service processor, system board, and I/O enclosure boards is released periodically. 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 4.2, “CPC failover and failback” on page 116.
6. Update the latest flash drive firmware regularly because enhancements for efficiency, performance, and reliability are released often. A flash drive firmware update is a concurrent process in the DS8000 series family. The instructions 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 DS8A00 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

The Automatic Code Management checkbox in the GUI under **Settings** → **System** → **Advanced** was removed in version 10.1. The feature is now managed through integrated version control and policy-based configuration. Automatic code updates and synchronization are handled via connected repository settings, while administrators can define automation rules in the System Configuration Policies section. This change simplifies settings and aligns with modern development workflows.

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, an HMC Code Image Server function is available. 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 11-2 on page 411 shows the relevant menu options in the Updates menu of the service console Web User Interface (WUI).

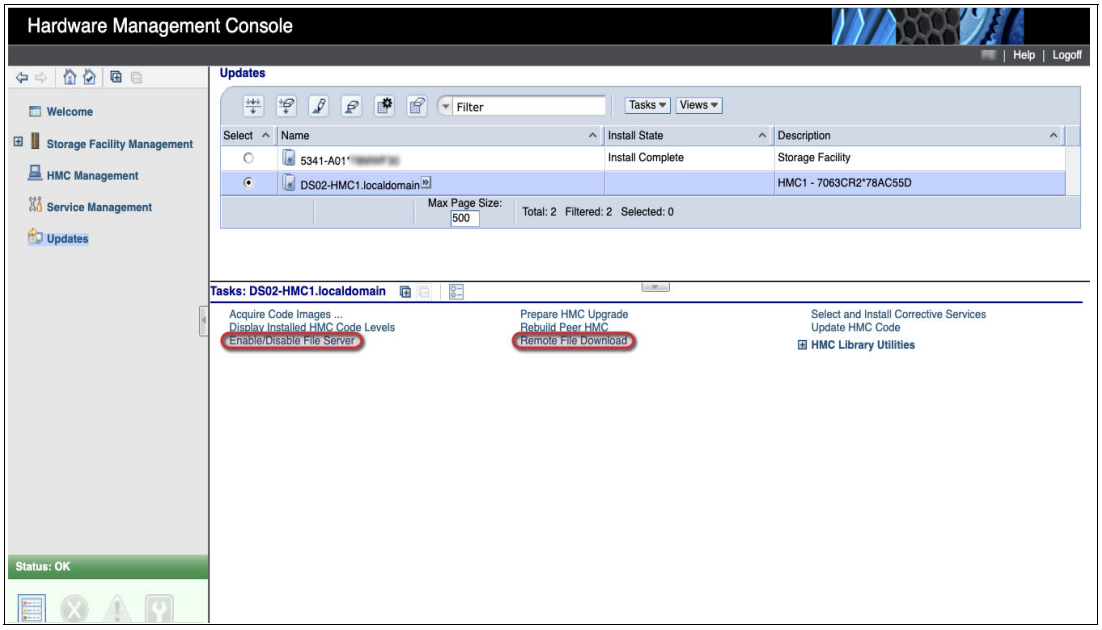


Figure 11-2 Updates menu: HMC Code Image Server

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 11-3.

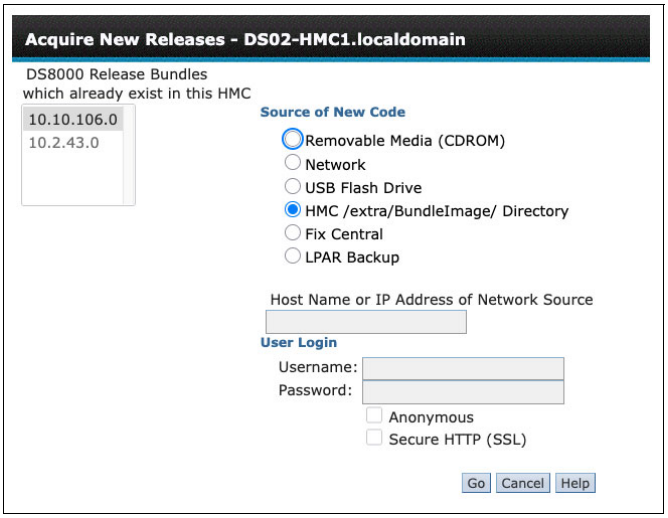


Figure 11-3 Acquiring from the HMC Directory

After the acquisition is complete, the normal code load process proceeds with the CDA software preinstallation.

Note: Bundles that were acquired from physical 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 11-3 on page 412 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.

11.2.2 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 DS8A00, 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 11-4.

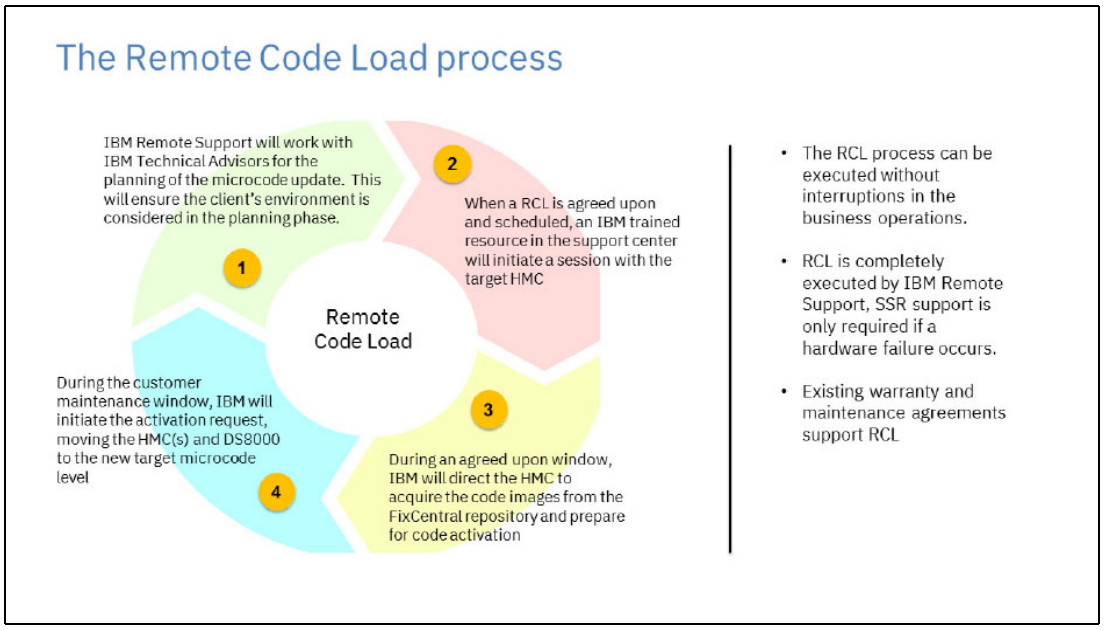


Figure 11-4 Remote Code Load process

The following steps are used for the Remote Code Load process:

1. IBM Remote Support Personnel work with IBM Technical Account Managers (TAM) or with the client directly to plan the microcode update.
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.

The RCL process, the prerequisites and other useful information is described in detail here:
https://www.ibm.com/support/pages/node/6590739

RCL appointments can be requested here:
https://ds8000.ibmrcl.enterpriseappointments.com/v2/#page/instructions

RCL requires access to the following servers (Table 11-1).

Table 11-1 Servers list

Host name	IPv4	Port(s)
esupport.ibm.com	192.148.6.11	443
esupport.ibm.com	170.225.123.67	443
www-945.ibm.com	192.148.6.11	443
www-945.ibm.com	170.225.123.67	443

Notes:

- Code bundles are pulled to the HMC. They are not pushed.
- IP addresses might change in the future. For this reason, it is strongly recommended to use host names for firewall configuration instead of IP addresses.

11.2.3 Customer Code Load

The IBM DS8000 G10 and the DS8900F at current code levels offer a feature that allows clients to upgrade their DS8000 storage systems independently. This option benefits organizations that prefer full control over planning, scheduling, and executing the code upgrade. The entire process is managed by the customer, without requiring IBM Remote Support personnel.

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.

This option is available under **Settings** → **Support** → **Update System**, as shown in Figure 11-5.

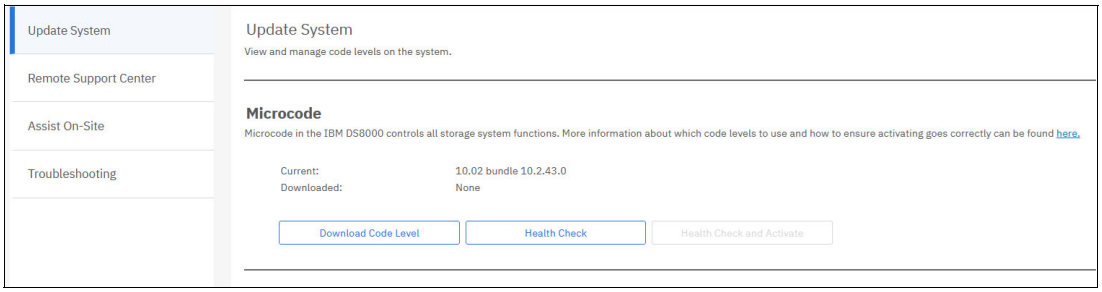


Figure 11-5 Storage Manager GUI: Update System menu

Figure 11-5 on page 414 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 11-5 on page 414, 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 11-6.

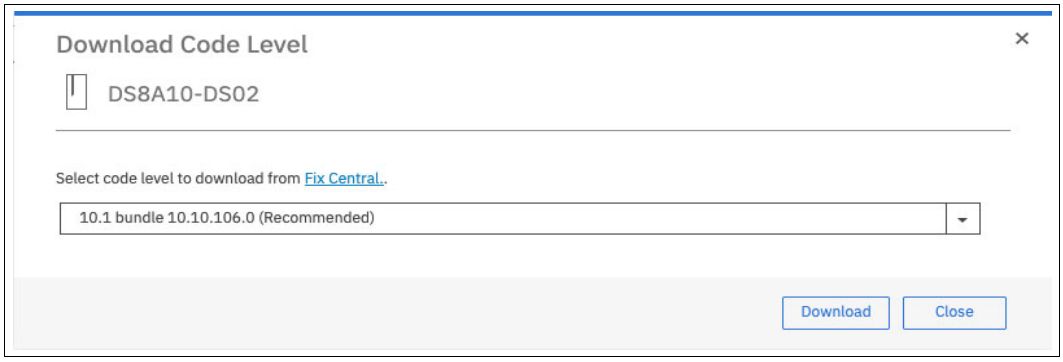


Figure 11-6 Microcode level query

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 11-7 on page 415.

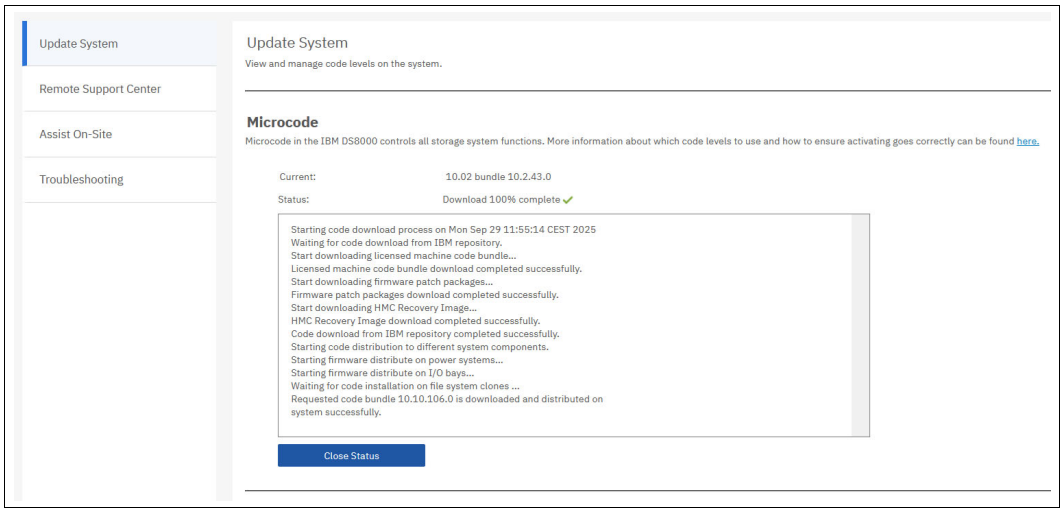


Figure 11-7 Microcode download complete message

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 11-8.

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 is still performed before and after the activation when the user selects **Health Check and Activate**.

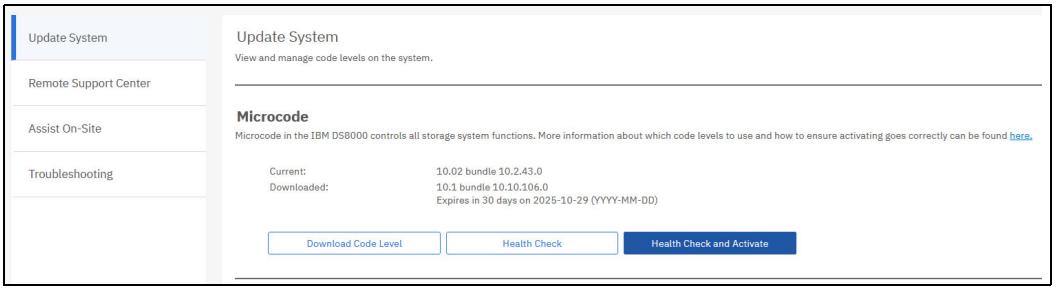


Figure 11-8 select Health Check and Activate

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 11-9 on page 416. To start, click **Yes**.

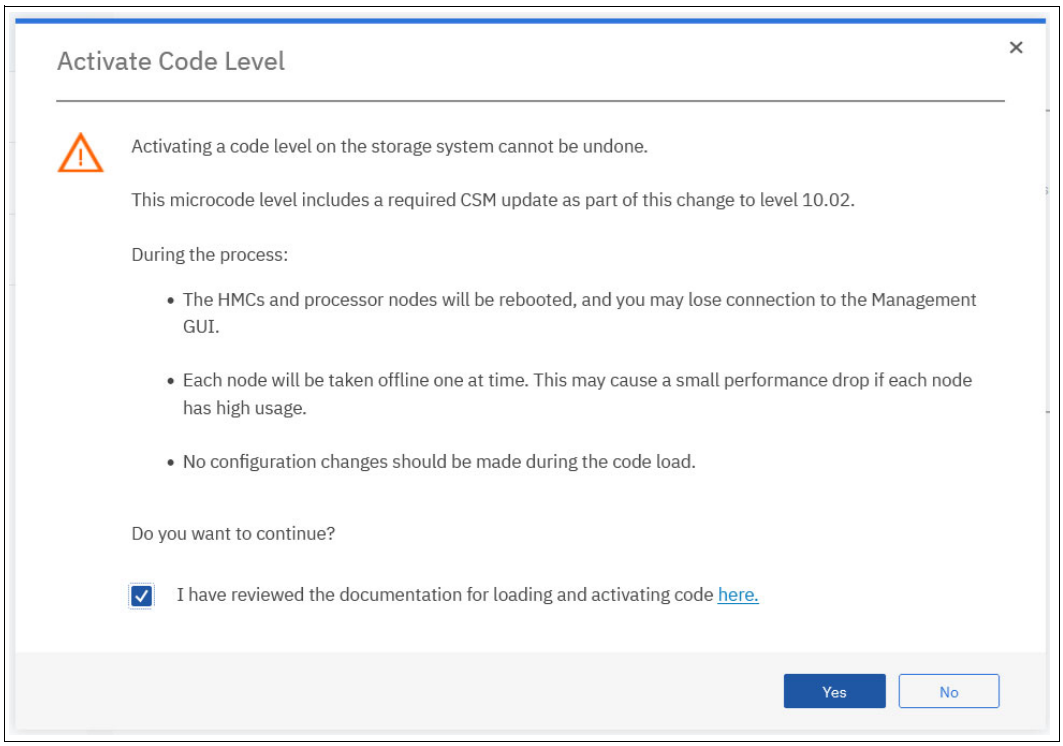


Figure 11-9 Code Activation Message

6. 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 11-10).

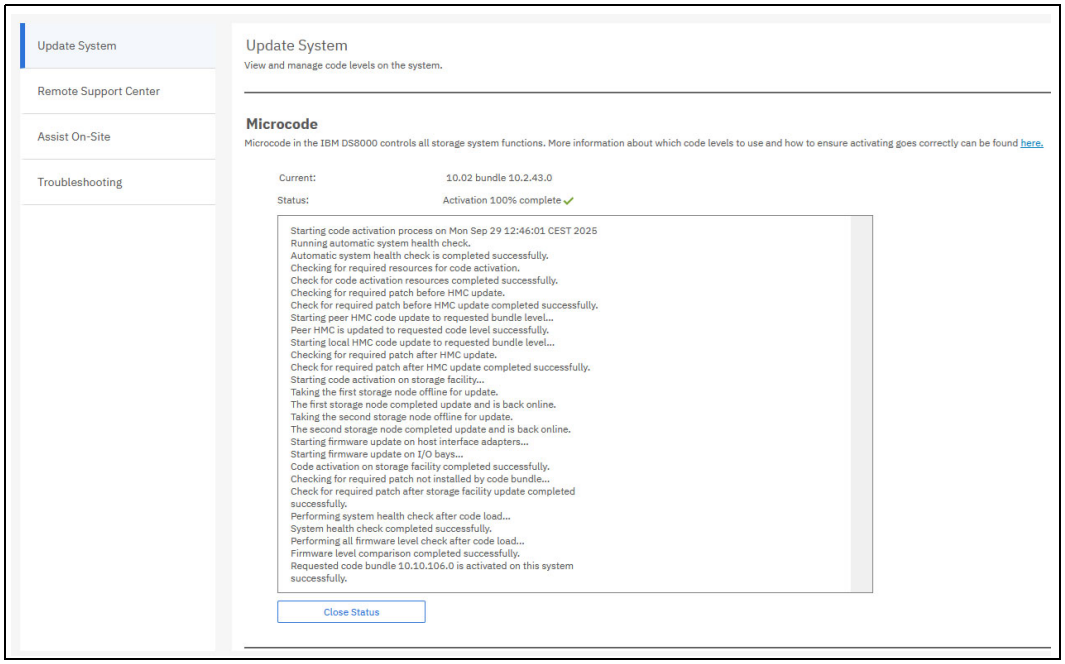


Figure 11-10 Code Activation Complete

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 is a case number in that notification, which can be used as a reference for when the client engages IBM Remote Support.

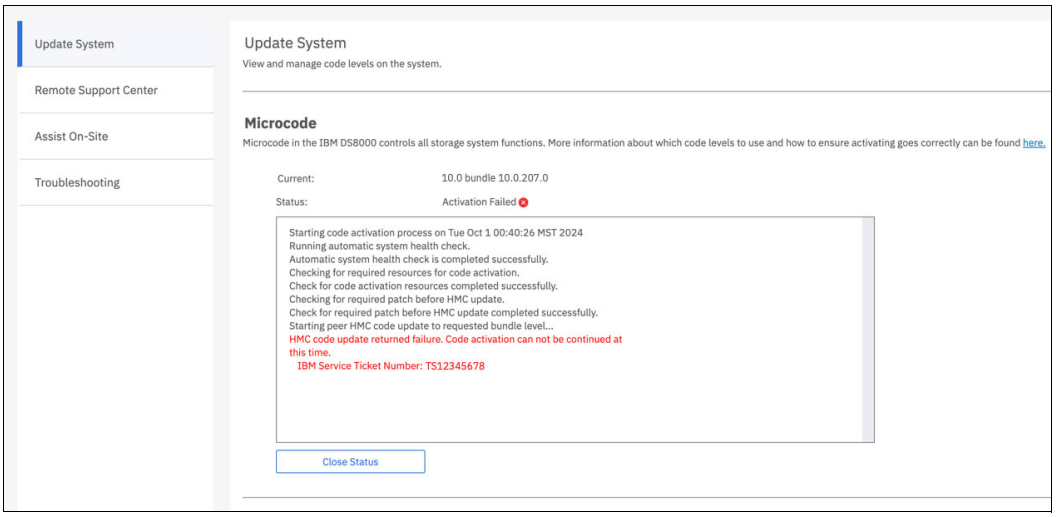


Figure 11-11 Case number

When the problem is resolved, the Code Load can be resumed by going back to step 5 on page 416 and select **Health Check and Activate.**

Note: If the option **Health Check and Activate** is not available, the Code Load needs to be restarted by going back to step 1 on page 414 to download the **same** Code Bundle which failed the activation.

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 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 or Technical Account Manager about proposed code load schedules.

Code recommendations are listed at DS8000 Code Recommendation.

Additionally if applicable, 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 DS8A00 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 can also 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 DS8A00 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 adapter firmware update because the DS8A00 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 DS8000 service actions. CUIR is available for the DS8000 when the DS8000 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 DS8000 service actions.

CUIR allows the DS8000 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 DS8000 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 DS8A00.

CUIR is disabled by default. You can enable it through the GUI by navigating to **Settings** → **Advanced**, as shown in Figure 11-12 on page 420. Enabling CUIR is a concurrent operation, meaning it does not require a system restart. However, the CUIR feature itself will only take effect during the next code load action.

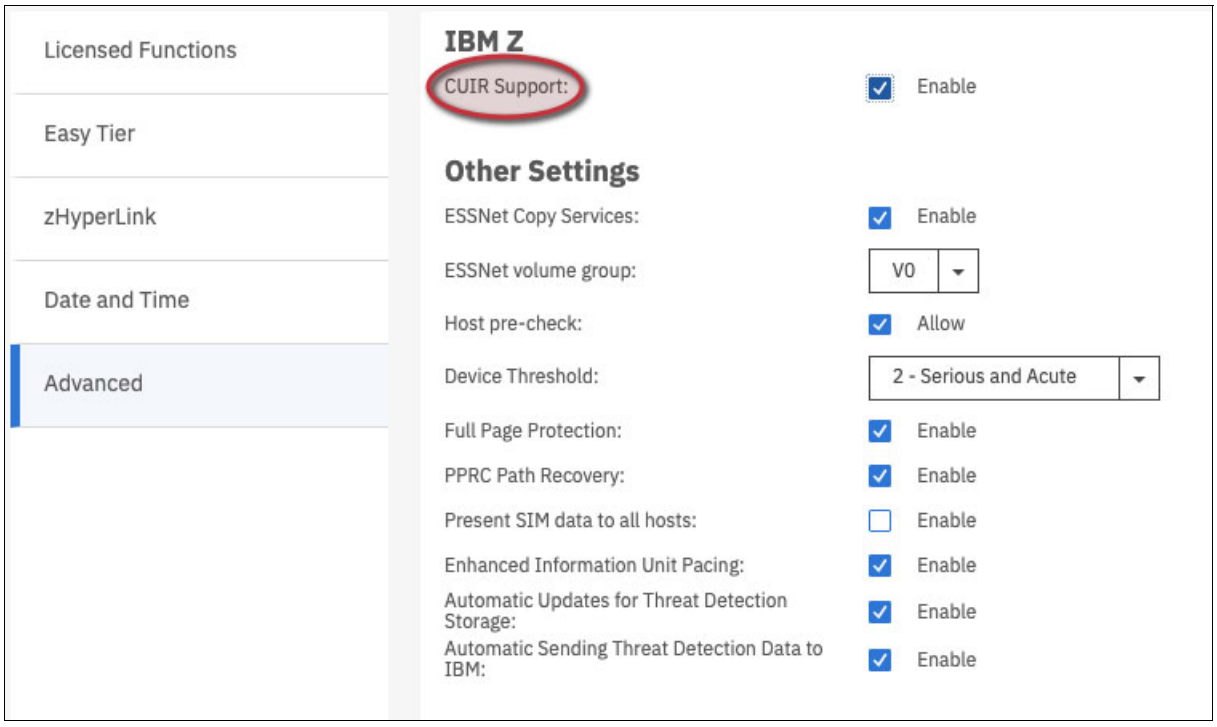


Figure 11-12 Settings - Advanced - CUIR Support

11.5 Loading the code bundle

The DS8000 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.

Loading the microcode can also 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 11.2.3, “Customer Code Load” on page 414.

The microcode can be loaded by an IBM SSR onsite as well. To review and arrange the required services, contact your IBM SSR or your IBM Technical Account Manager.

The CCL process includes an *autonomic recovery function*, which means that it is tolerant to temporary non-critical errors that might surface during the activation. During the CCL, 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 DS8000 system can continue with the code update with tolerable errors. 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 so that DS8A00 clients can 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. DS8000 development always strives to reduce the firmware activation time to improve system redundancy 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 DS8000 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.

11.6 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 DS8000.
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 DS8000.
5. Verify the DS Storage Manager connectivity from IBM Copy Services Manager to the DS8000.
6. Verify the connectivity from the DS8000 to all IBM Security Guardium Key Lifecycle Manager instances, or other servers in use.

11.7 Summary

IBM might release changes to the DS8000 LMC. These changes might include code fixes and feature updates that relate to the DS8000.

These updates and the information about them are documented in the DS8A00 Code cross-reference website. You can find this information for a specific bundle under the Bundle Release Note information section in DS8000 Code Recommendation.

12

Monitoring and support

This chapter provides information about the Simple Network Management Protocol (SNMP) implementation and messages for the IBM DS8A00 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 DS8000” on page 424
- “SNMP notifications” on page 425
- “SNMP configuration” on page 433
- “Introducing remote support” on page 437
- “IBM policies for remote support” on page 438
- “Remote support advantages” on page 438
- “Remote support and Call Home” on page 438
- “Remote Support Access (inbound)” on page 440
- “Call Home and Assist On-Site customer-provided certificates” on page 445
- “RESTful API” on page 447
- “Audit logging” on page 450
- “Using IBM Storage Insights” on page 451
- “IBM Call Home Connect Cloud” on page 459

© Copyright IBM Corp. 2025. All rights reserved.

12.1 SNMP implementation on the DS8000

SNMP, as used by the DS8000, is designed so that the DS8000 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 DS8A00 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 DS8A00, 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 DS8000 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
coldStart	0	Restart after a crash.
warmStart	1	Planned restart.
linkDown	2	Communication link is down.
linkUp	3	Communication link is up.
authenticationFailure	4	Invalid SNMP community string was used.
egpNeighborLoss	5	Exterior Gateway Protocol (EGP) neighbor is down.
enterpriseSpecific	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 DS8A00 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 DS8A00 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 DS8A00
- 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 DS8A00

```
Manufacturer=IBM
ReportingMTMS=2107-A05*78NFT51
ProbNm=3084
LparName=SF78NFT50ESS01
FailingEnclosureMTMS=2107-A05*78NFT51
```

```
SRC=BE80CB13
EventText=Recovery error,the device hardware error.
FruLoc=Part Number 02AT314 FRU CCIN Y7G1
FruLoc=Serial Number 8CE38EE300A19242
FruLoc=Location Code U2107.D05.UL000GC-P3-C6
```

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 1*xx* traps are sent for a state change of a physical link connection. The 2*xx* 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 DS8000. For more information about these functions and terms, see *IBM DS8000 Copy Services,* SG24-8367.

Physical connection events

Within the trap 1*xx* 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-A08 78LRV51 12
SEC:  IBM 2107-A05 78NFT51 24
Path: Type  PP   PLink  SP   SLink  RC
1:    FIBRE 0143 XXXXXX 0010 XXXXXX 15
2:    FIBRE 0213 XXXXXX 0140 XXXXXX OK
```

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-A08 78-LRV51 10
SEC:  IBM 2107-A05 78-NFT51 20
Path: Type  PP   PLink  SP   SLink  RC
1:    FIBRE 0143 XXXXXX 0010 XXXXXX 17
2:    FIBRE 0213 XXXXXX 0140 XXXXXX 17
```

After the DS8000 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-A08 78-LRV51 21
SEC:  IBM 2107-A05 78-NFT51 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-A08 78-LRV51 84 08
SEC:  IBM 2107-A05 78-NFT51 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 Suspended 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-A08 78-LRV51 28 00 01
SEC:  IBM 2107-A05 78-NFT51 a8 00
Start: 2025/11/14 10:30:32 CET
PRI Dev Flags (1 bit/Dev, 1=Suspended):
C0000000000000000000000000000000000000000000000000000000000000000
```

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-A08 78-LRV51
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-A05 78-NFT51
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-A05 78-NFT51
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-A05 78-NFT51
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-A05 78-NFT51
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-A08 78-LRV51
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-A05 78-NFT51
Subordinate:  IBM 2107-A05 78-HAL91
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-A05 78-NFT01
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-A05 78-NFT51
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-A05 78-NFT51
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-A05 78-NFT51
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-A05 78-NFT01
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-A05 78-NFT01
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 DS8A00 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-A05 78-NFT51
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-A05 78-NFT51
Extent Pool ID: P1
Limit: 95%
Threshold: 95% Status: 0
```

12.2.4 Compression SNMP

With DS8A00, a new specific trap is added that relates to compression for FCMs. This trap is sent out when the compression capacity of an extent pool goes above the specified threshold, as shown in Example 12-22

The following information is shown for specific trap 228:

- ssEventDescr=yyyy/MM/dd hh:mm:ss tzn

  Shows the time that the alert was sent (year, month, day, hour, minute, second, and time zone)

- -error_description-

  Shows the event description

- UNIT: Mnf Type-Mod SerialNm
- HMC IP: -ipaddr-
  - shows the IPv4 or IPv6 address of the HMC sending the trap
- Extent Pool ID: ee

  ee is the extent pool identifier, for example: P0

- Compression Threshold: kk%

  kk is the compression capacity threshold percentage

- Compression Status: ss

  The following statuses can be seen:
  - Above Compression Threshold
  - Pool Physical Capacity Full
  - Pool Physical Capacity Full, Emergency Reserve 1
  - Pool Physical Capacity Full, Emergency Reserve 2

Example 12-22 Compression SNMP

```
Frame Length: 315 bytes
Version: SNMPv1
Community: public
OID: .1.3.6.1.4.1.2.6.130
Address: 9.11.222.162
sysUpTime: 0 days, 00:36:47
Generic: 6 - Enterprise Specific
Specific: 228
OID: .1.3.6.1.4.1.2.6.130.3.2
ASN1 Type: Octet String 0x04 (4)
Value: 2025/10/09 14:28:31 CEST
Extent Pool Compression Capacity Above Treshold
UNIT: Mnf Type-Mod SerialNm
      IBM 2107-A05 13-02211
HMC IP: 9.11.222.162
Extent Pool ID: P12
Compression Threshold:  15%
Compression Status: Above Compression Threshold
```

12.3 SNMP configuration

The SNMP for the DS8000 is designed to send traps as notifications. The DS8000 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 DS8A00 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 Alert traps section of the IBM Storage DS8000 (10th Generation) 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 DS8A00 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 initial 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 437.

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.

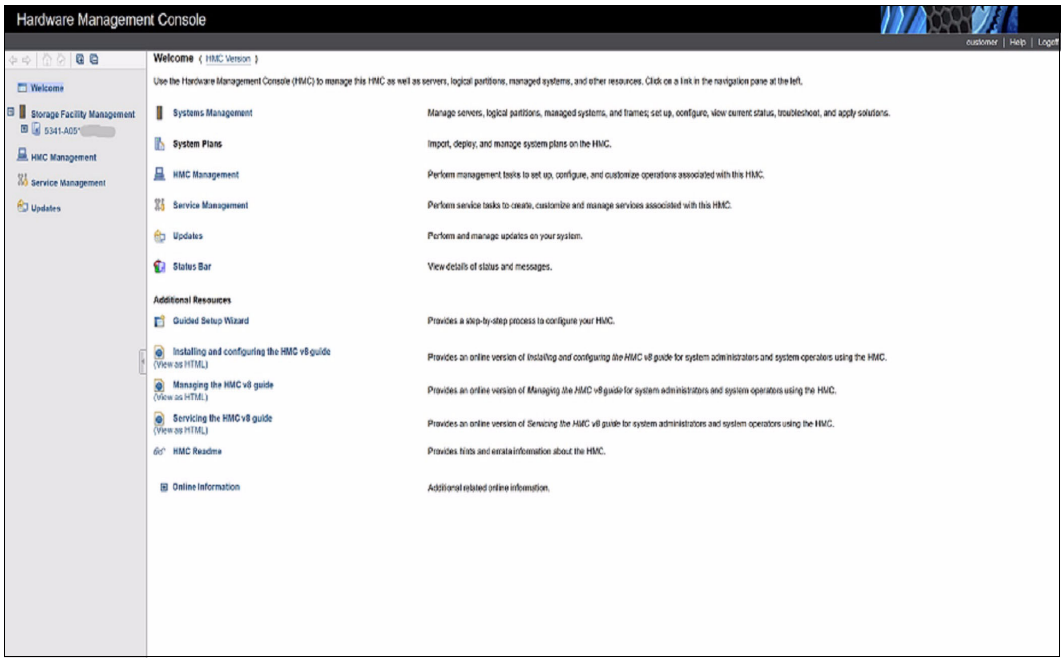


Figure 12-1 HMC Service Management

2. Select **Manage Serviceable Event Notification**, as shown in Figure 12-2.

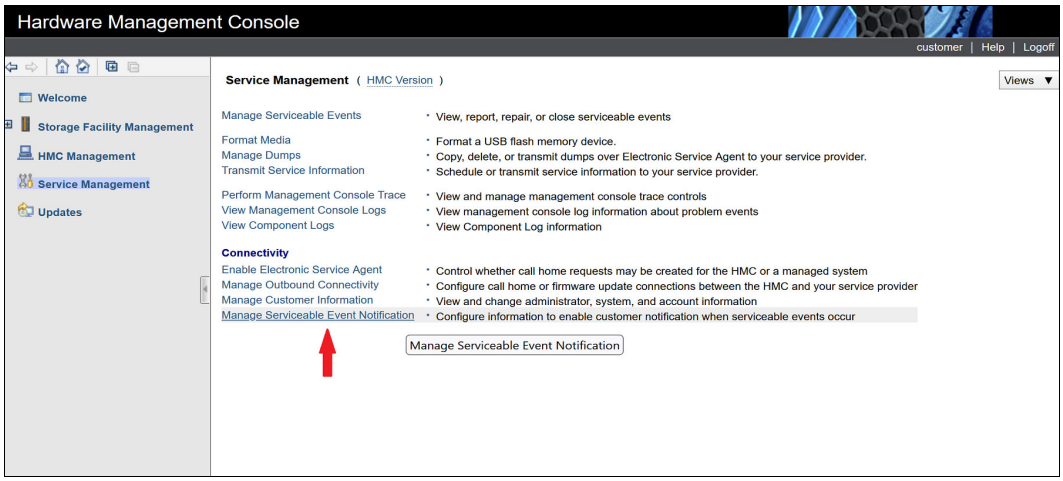


Figure 12-2 HMC Manage Serviceable Event Notification Selection

3. A new window opens. Enter the TCP/IP information of the SNMP server in the SNMP Trap Configuration dialog. Select Trap 3, as shown in Figure 12-3.

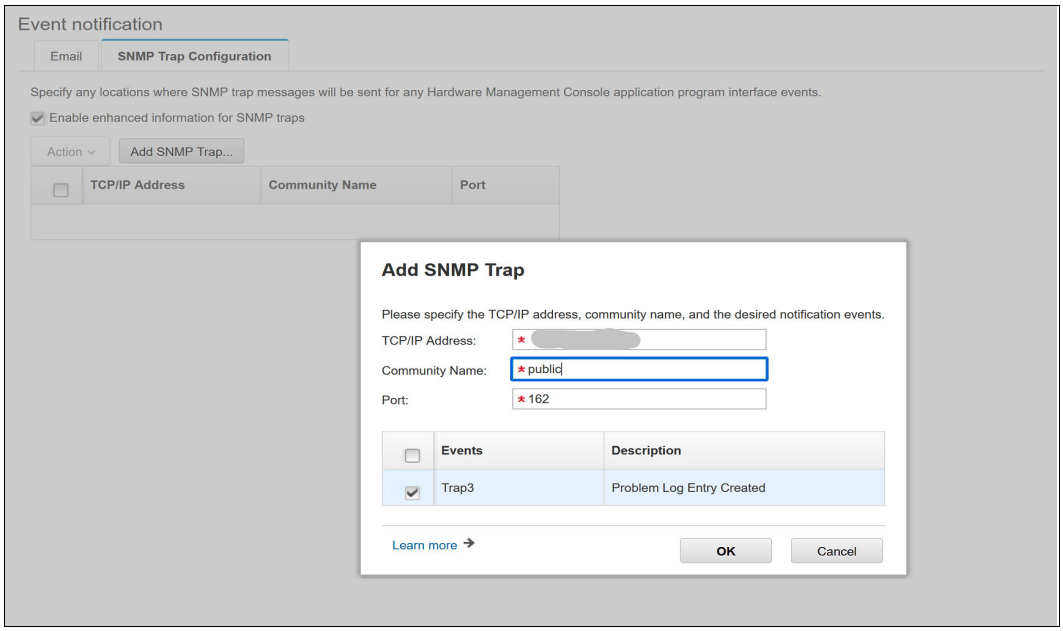


Figure 12-3 MC Manage Serviceable Event Notification

4. To verify the successful setup of your environment, create a Test Event on your DS8A00 MC by selecting the IP address and **Test SNMP Trap**, as shown in Figure 12-4.

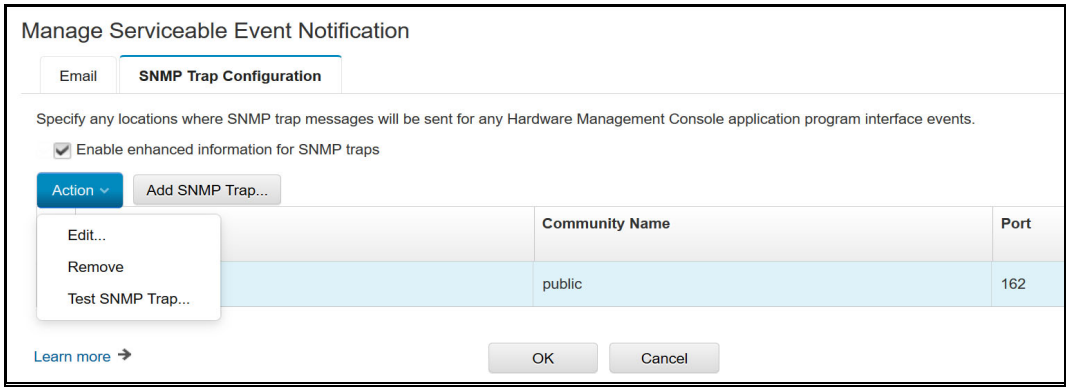


Figure 12-4 HMC test SNMP trap

Check the SNMP server for the successful reception of the test trap.

5. The test generates the Service Reference Code BEB20010, and the SNMP server receives the SNMP trap notification, as shown in Figure 12-5 on page 436.

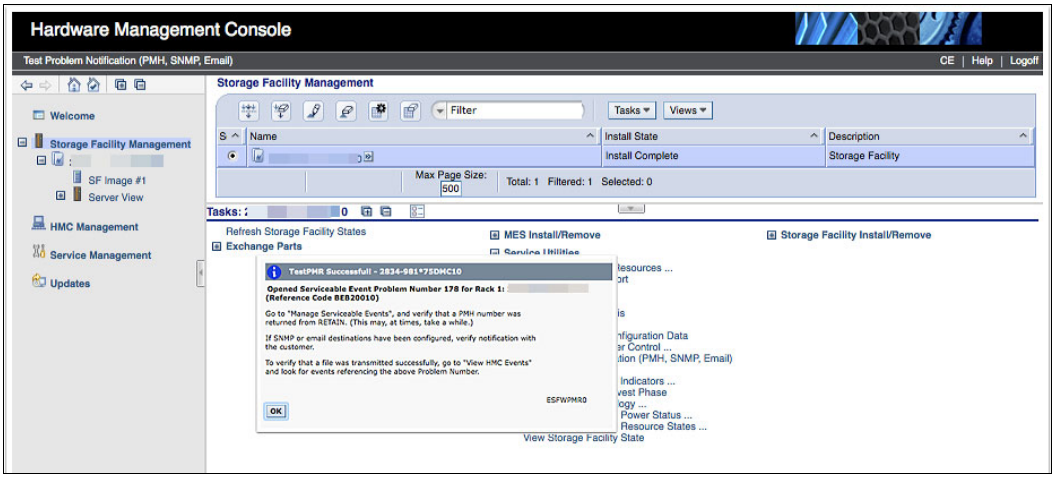


Figure 12-5 HMC SNMP trap test

Errors occurring after configuring the SNMP traps on the HMC are sent to the SNMP server, as shown in Example 12-23.

Example 12-23 Example of an error that is reported to the SNMP system

```
2025-09-09 17:20:19 srv22-755.dcry.i [10.7.55.22] (via UDP:
[10.7.55.22]:53254->[0.0.0.0]:0) TRAP, SNMP v1, community public
   Ibm2100-MIB::ibm2100 Enterprise Specific Trap (Ibm2100-MIB::ssSystemNotifyTrap)
Uptime: 17 days, 23:23:47.05
   Ibm2100-MIB::ssEvent = STRING: "Sep 9, 2025 5:18:59 PM CEST
Manufacturer=IBM
ReportingMTMS=5341-A05*78HAL90
ProbNm=341
LparName=SF78HAL90ESS11
FailingEnclosureMTMS=5341-A05*78HAL90
SRC=BE340040
EventText=Health check found fenced adapter.
Fru1Loc=U1700.1B4.RJB00HT-P1-C2
Fru2Loc=n/a

PMH=TS020251650
Reporting HMC Hostname=DS02-HMC1.localdomain."
```

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 12-24 shows how SNMP is enabled by using the **chsp** command.

Example 12-24 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
SNMPadd         10.10.10.1,10.10.10.2
emailnotify     Disabled
emailaddr       -
emailrelay      Disabled
emailrelayaddr  -
emailrelayhost  -
numkssupported  64
loginprompt     Off
promptmessage   -
guisuggestions  show
```

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 DS8A00 specific Message Information Base file.

The Message Information Base file that is delivered with the latest DS8A00 DS CLI CD is compatible with all previous levels of DS8A00 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 DS8A00 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 DS8A00 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 DS8A00:

- 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 DS8000 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 DS8A00 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 working is sent to an IBM SSR. The DS8A00 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 DS8000 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 DS8A00 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 DS8000 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 secure 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 Storage DS8000 (10th Generation) Introduction and Planning Guide, G10-I53-01.

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 12-25 shows exporting the configuration files.

Note: The **offloadfile** command cannot be run from the embedded DS CLI window.

Example 12-25 DS CLI command to download configuration files

```
dscli> offloadfile -config c:\dscli_offload
Date/Time: 09 October 2025 13:48:58 CEST IBM DSCLI Version: 7.10.10.348 DS:
IBM.2107-78HAL91
CMUC00428I offloadfile: The config file has been offloaded to
c:\dscli_offload\productSettings.
CMUC00428I offloadfile: The config file has been offloaded to
c:\dscli_offload\ICSInstall.history.
```

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 DS8A00, 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 1.3, but TLS 1.2 is also supported. 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 DS8000 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 442.

If 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 DS8000 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 DS8000 systems or other IBM storage systems.

A simple AOS connection to the DS8000 is shown in Figure 12-6. For more information about AOS, prerequisites, and installation, see *IBM Assist On-site for Storage Overview*, REDP-4889.

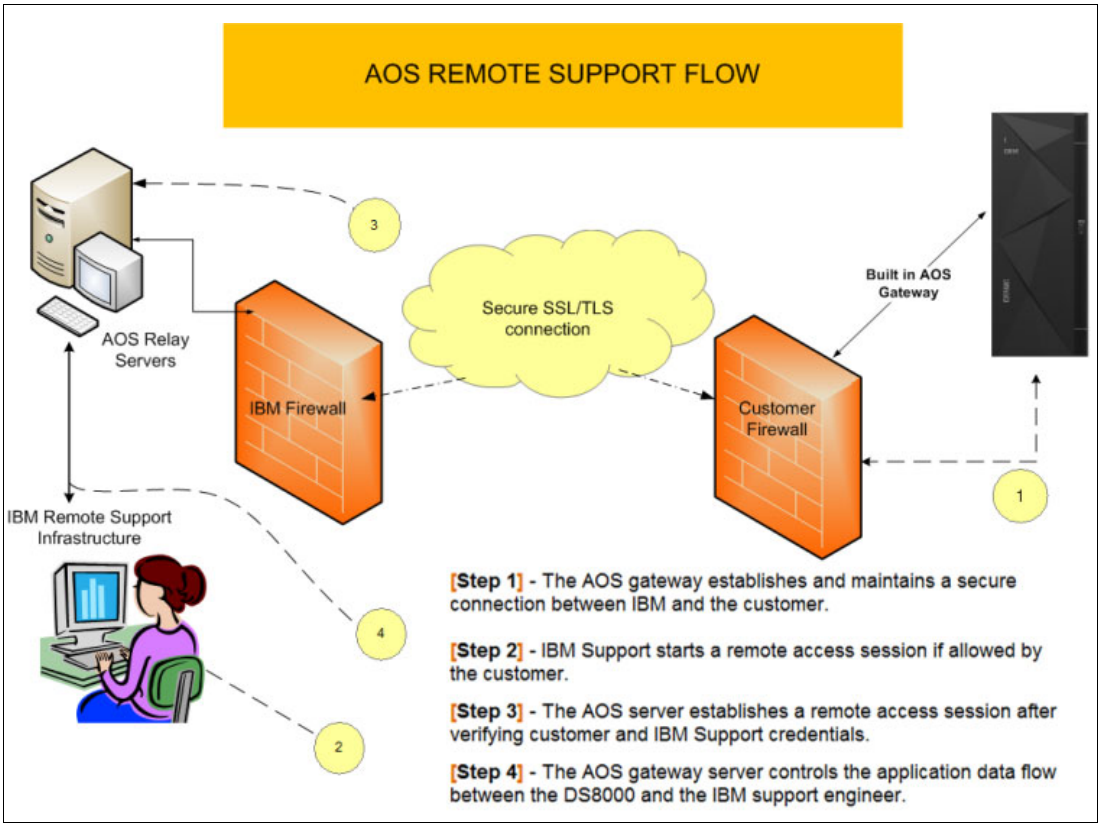


Figure 12-6 DS8000 AOS connection

12.8.2 DS8A00-embedded AOS

AOS is an embedded feature on DS8000 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 Storage DS8000 (10th Generation) Introduction and Planning Guide, G10-I53-01.

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 DS8000

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 12-7.

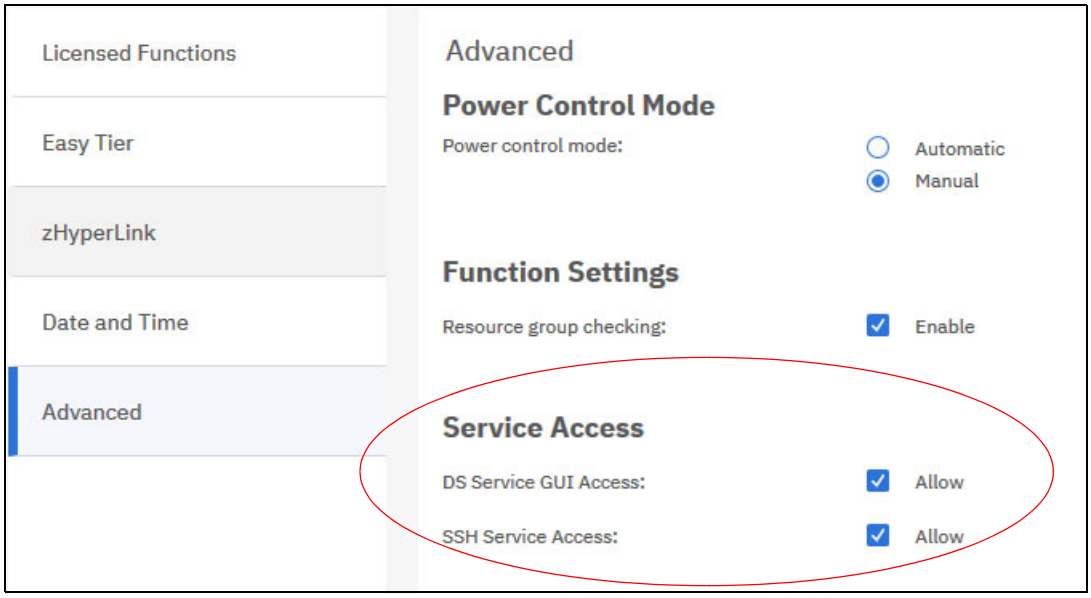


Figure 12-7 Controlling service access through the DS Storage Manager GUI

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 12-8.

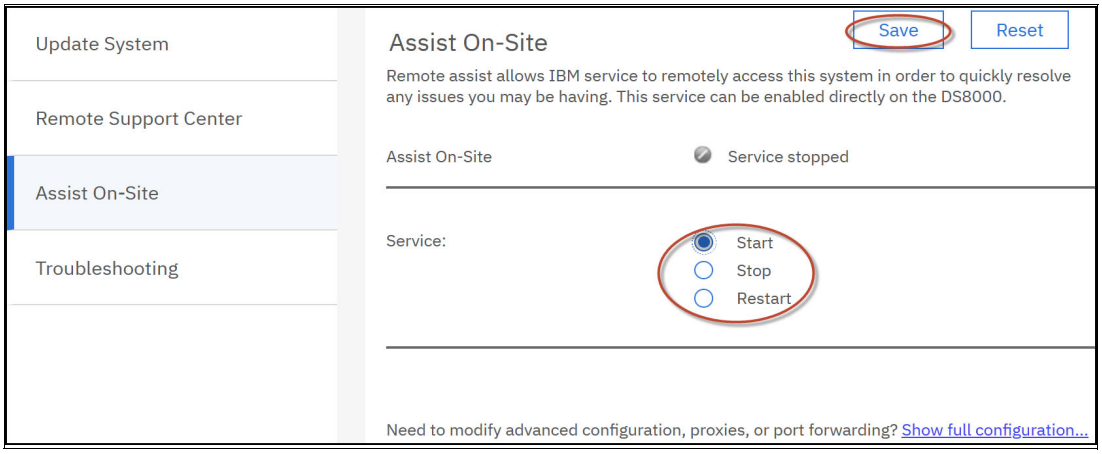


Figure 12-8 Managing the Assist On-Site connection

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. Provide the code to the IBM Remote Support Personnel that attempt to connect to the storage system. Figure 12-9 shows the Remote Support Center window.

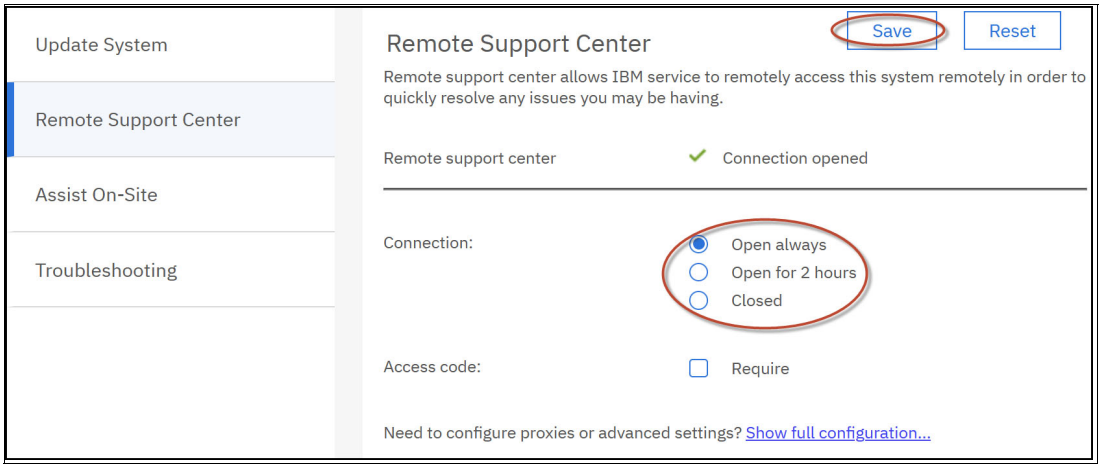


Figure 12-9 Managing the Remote Support Center connection

Using DS CLI to manage service access

You can manage service access to the DS8A00 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 12-26.

  ```
  chaccess [-commandline enable|disable] [-wui enable|disable] [-aos enable|
  disable] [-rsc enable|disable] [-hmc 1|2|all]
  ```

  Example 12-26 Output of the chaccess command

  ```
  dscli> chaccess -aos enable -hmc 1
  CMUC00441I chaccess: hmc1: The access settings were successfully modified.
  ```

- **lsaccess**: This command displays the access settings of the primary and backup MCs:

  ```
  lsaccess [-hmc 1|2|all]
  ```

  See the output in Example 12-27.

  Example 12-27 lsaccess command output

  ```
  dscli> lsaccess -hmc all -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
  ```

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 can also be configured at the MC to send notifications in a remote support connection.

12.9 Call Home and Assist On-Site customer-provided certificates

Figure 12-10 shows an overview of the communication path for Call Home with customer proxy configured and using a certificate that is provided by the customer.

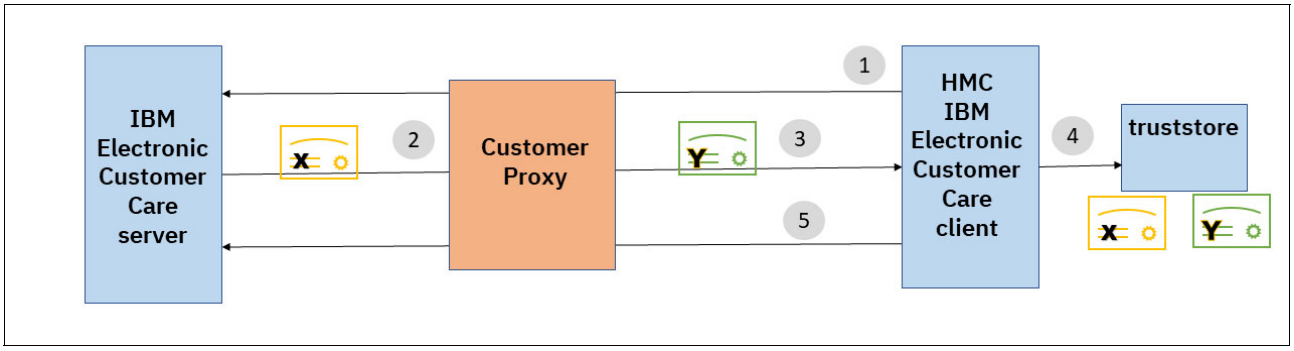


Figure 12-10 Call Home customer-provided certificate - Communication path

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 12-11 shows an overview of the communication path for AOS with customer proxy configured and using a certificate that is provided by the customer.

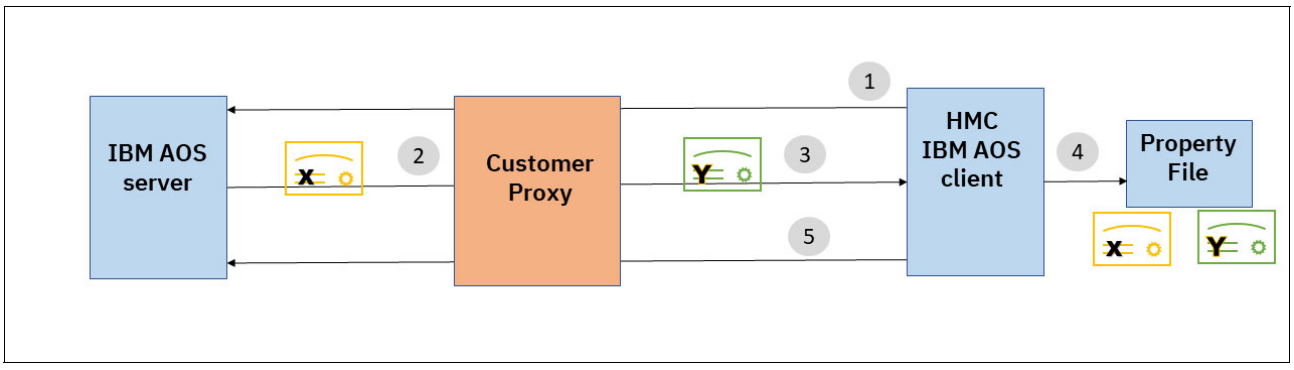


Figure 12-11 AOS customer-provided certificate - Communication path

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 12-12 shows these settings in the DS GUI. For detailed configuration steps, see `Call home settings`.

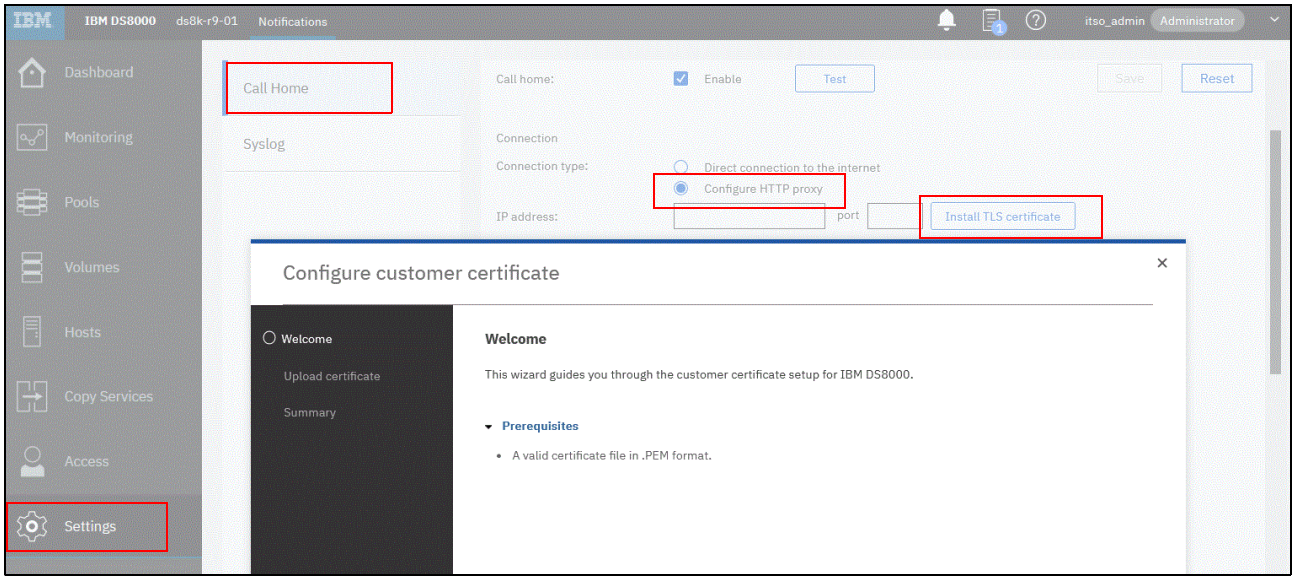


Figure 12-12 Installing a customer-provided certificate for Call Home

To install a customer-provided certificate for the AOS connection, from the DS GUI access **Settings** → **Support** → **Assist On-Site** → select **Configure HTTP proxy** → click **Install TLS certificate**. If initially the setting is not visible, click **Show full configuration** at the bottom of the page. Figure 12-13 shows these settings in the DS GUI. For detailed configuration steps, see Assist On-Site.

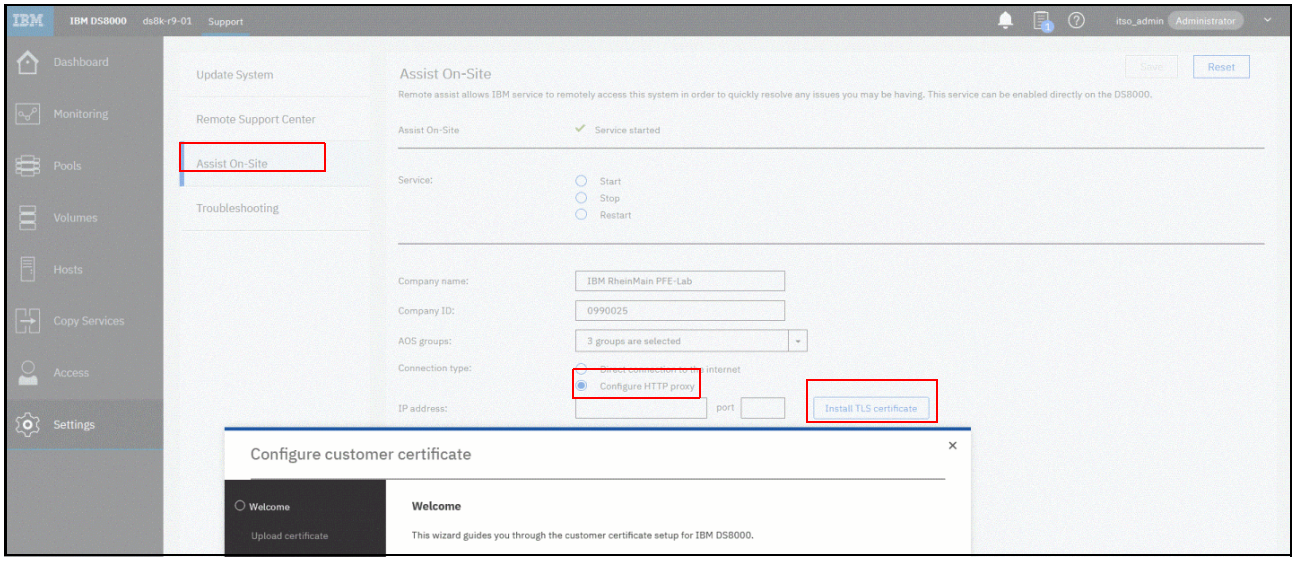


Figure 12-13 Installing a customer-provided certificate for AOS

Guidance from IBM support is available to help setting up this configuration.

12.10 RESTful API

The RESTful API is an application on your storage system HMC for initiating simple storage operations through the web.

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

The RESTful API is intended for development, testing, and debugging of the client management infrastructures in your storage system. You can use the RESTful API with a CURL command or through standard web browsers.

A command reference for available REST operations is provided in the RESTful API documentation.

You can also consult the *RESTful API Guide*, SC27-9823, or *Exploring the DS8870 RESTful API Implementation*, REDP-5187.

Also, the IBM Copy Services Manager provides a REST API, as documented at this link.

Example 12-28 shows how to start with a REST session against the DS8000. As a first action, you post your username and password to the HMC in exchange for a session token.

Example 12-28 Starting with REST: Getting token

```
C:\Program Files\IBM\dscli>curl -k -H "Content-Type: application/json" -X POST -d
"{request:{params:{username:itso_admin,password:Test1234}}}"
https://10.7.55.22/api/v1/tokens | python -m json.tool
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100   264  100   202  100    62    159     48  0:00:01  0:00:01 --:--:--   208
{
    "server": {
        "status": "ok",
        "code": "",
        "message": "Operation done successfully."
    },
    "token": {
        "token": "9bd12124ac0e44f6a9a39db205b72e25",
        "expired_time": "2025-10-09T16:37:19+0200",
        "max_idle_interval": "1800000"
    }
}
```

Using the above token, you now can query various parts of the machine, for instance starting a query on overall DS system level, as shown in Example 12-29.

Example 12-29 Query the DS8000 on system level

```
C:\Program Files\IBM\dscli>curl -k -H
"X-Auth-Token:9bd12124ac0e44f6a9a39db205b72e25" -X GET
https://10.7.55.22/api/v1/systems | python -m json.tool
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100   407  100   407    0     0     61      0  0:00:06  0:00:06 --:--:--    96
```

```
{
    "server": {
        "status": "ok",
        "code": "",
        "message": "Operation done successfully."
    },
    "counts": {
        "data_counts": 1,
        "total_counts": 1
    },
    "data": {
        "systems": [
            {
                "id": "2107-78HAL91",
                "name": "DS8A10-DS02",
                "state": "online",
                "release": "10.1",
                "bundle": "10.10.106.0",
                "MTM": "5341-A01",
                "sn": "78HAL91",
                "wwnn": "5005076303FFF1D5",
                "cap": "256317936000000",
                "capalloc": "28460064000000",
                "capavail": "227850752000000",
                "capraw": "163840000000000"
            }
        ]
    }
}
```

Likewise, it is possible to query machine on other levels, for instance the extent pools, as in Example 12-30.

Example 12-30 Query the DS8000 on Extent Pool level

```
C:\Program Files\IBM\dscli>curl -k -H
"X-Auth-Token:9bd12124ac0e44f6a9a39db205b72e25" -X GET
https://10.7.55.22/api/v1/pools | python -m json.tool
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100  3019  100  3019    0     0    236      0  0:00:12  0:00:12 --:--:--   853
{
    "server": {
        "status": "ok",
        "code": "",
        "message": "Operation done successfully."
    },
    "counts": {
        "data_counts": 6,
        "total_counts": 6
    },
    "data": {
        "pools": [
            {
                "id": "P0",
                "link": {
```

```
                    "rel": "self",
                    "href": "https://10.7.55.22/api/v1/pools/P0"
                },
                "name": "CKD_Boot",
                "node": "0",
                "stgtype": "ckd",
                "cap": "12982158",
                "capalloc": "6554100",
                "capavail": "6426945",
                "overprovisioned": "1.7",
                "easytier": "managed",
                "tieralloc": [],
                "threshold": "15",
                "real_capacity_allocated_on_ese": "336",
                "virtual_capacity_allocated_on_ese": "16370008",
                "extent_size": "21cyl",
                "eserep": {},
                "tserep": {},
                "volumes": {
                    "link": {
                        "rel": "self",
                        "href": "https://10.7.55.22/api/v1/pools/P0/volumes"
                    }
                }
            },
            },
            ......
```

It is also possible to query for events of a certain category, for instance `Warning` events, as in Example 12-31.

Example 12-31 Query for events of a certain category

```
C:\Program Files\IBM\dscli>curl -k -H
"X-Auth-Token:9bd12124ac0e44f6a9a39db205b72e25" -X GET
https://10.7.55.22/api/v1/events?severity=warning | python -m json.tool
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100  3393  100  3393    0     0   4694      0 --:--:-- --:--:-- --:--:--  4692
{
    "server": {
        "status": "ok",
        "code": "",
        "message": "Operation done successfully."
    },
    "counts": {
        "data_counts": 1,
        "total_counts": 1
    },
    "data": {
        "events": [
            {
                "id": "SE88539",
                "type": "PoolDataReductionStatusModified",
                "severity": "warning",
                "time": "2025-07-02T14:48:10+0200",
```

```
                "resource_id": "P4",
                "formatted_parameter": [
                    "Mexico",
                    "Below Data Reduction Threshold",
                    "Pool Physical Capacity is Full"
                ],
                "description": ""
            }
        ]
    }
}
```


Other useful queries include those for FlashCopy or PPRC relationships and their status. Additionally, various POST, PUT, and DELETE requests can be used to update your DS8000 system.

12.11 Audit logging

The DS8000 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 12-32.

Example 12-32 DS CLI command to download audit logs

```
dscli> offloadauditlog -logaddr smc1 c:\temp\audit.txt
CMUC00243I offloadauditlog: Audit log was successfully offloaded from smc1 to
c:\temp\audit.txt.
```

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 12-14.

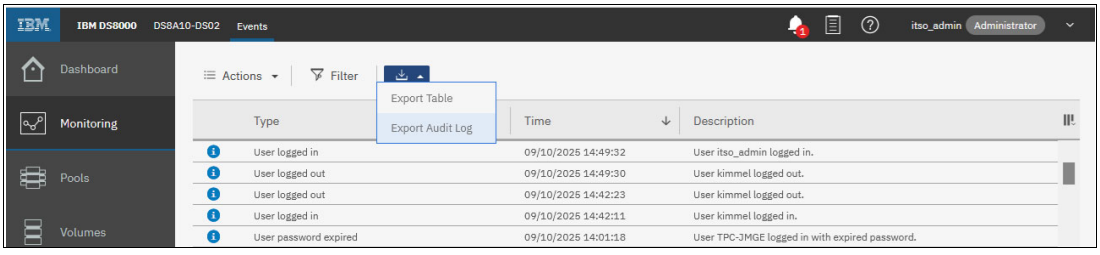


Figure 12-14 Export Audit Log

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 12-33.

Example 12-33 Audit log entries that relate to a remote support event

```
U,2025/11/25 12:22:38:000 CET,customer,1,IBM.2107-78HAL90,N,8020,
WUI_session_started,,,,
U,2025/11/25 12:40:36:000 CET,customer,1,IBM.2107-78HAL90,N,8022,
WUI_session_logoff,WUI_session_ended_loggedoff,,,
U,2025/11/26 11:48:33:000 CET,,1,IBM.2107-78HAL90,N,8036,Authority_to_root,Challen
ge Key = 'ABCD-EFGH-IJKL-1234'; Authority_upgrade_to_root,,,
```

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 DS8A00. 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 DS8000.

Entries are added to the audit file only after an operation completes, ensuring that all details about the request and its completion status are known. Each entry logs both the request and the response information. In rare cases, an operation may not complete due to a timeout; if this occurs, no entry is recorded 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.

12.12 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 DS8000.

12.12.1 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 DS8A00 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.

12.12.2 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 (bottom of that page) 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 12-15:

   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. At least 1 GB of RAM and 3 GB of disk space are required.

   c. Run **`installDataCollectorService.sh`** (Linux or AIX) or **`installDataCollectorService.bat`** (Windows).

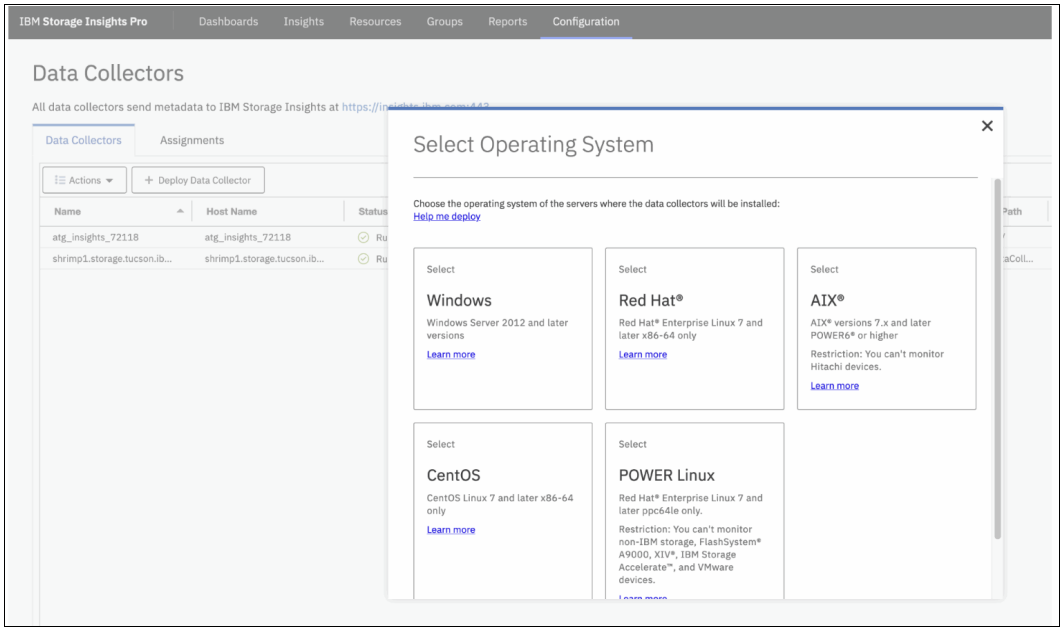


Figure 12-15 Deploy Data Collector

Figure 12-16 on page 454 shows a system overview where you see alerts, can access tickets and the actions that you can take to manage tickets.

Note: If you do not have an IBMid yet, go to IBM and register.

3. 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.

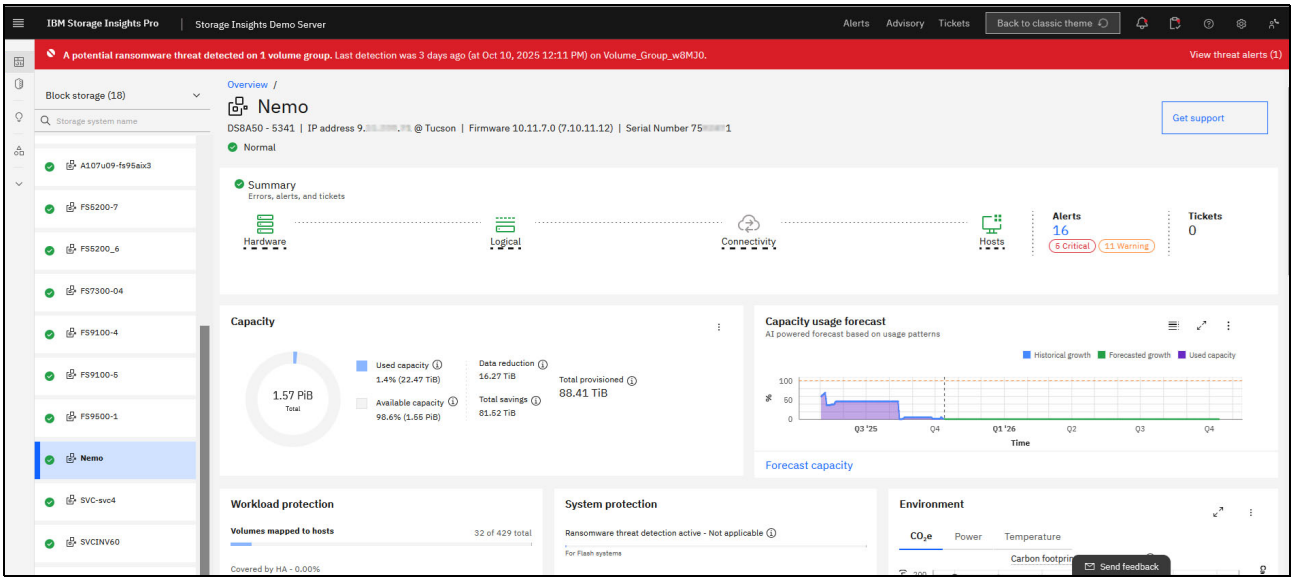


Figure 12-16 IBM Storage Insights: System overview

An example of an IBM Storage Insights Pro resource view is the detailed volume information shown in Figure 12-17.

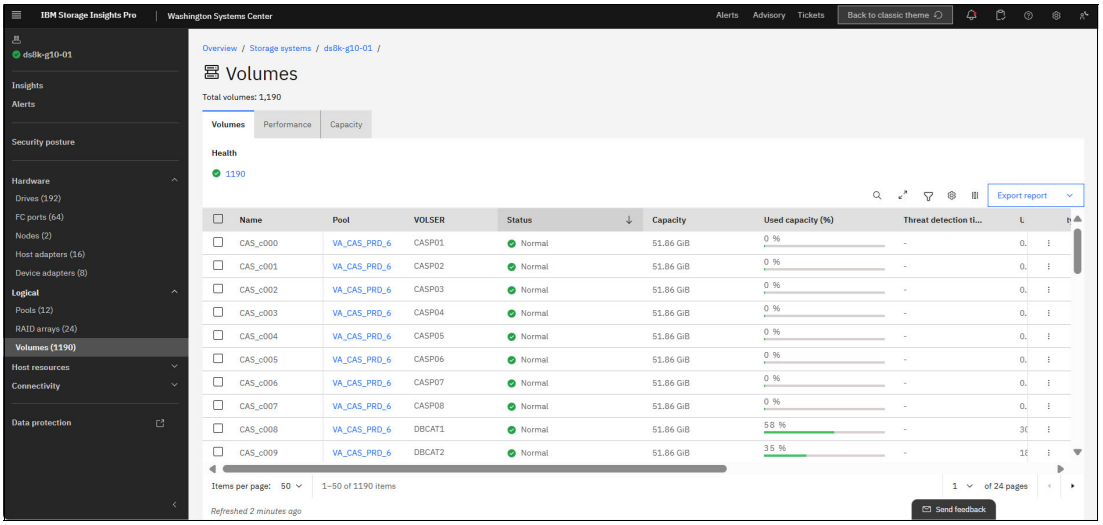


Figure 12-17 IBM Spectrum Insights Pro volume details

For more information to enable a proper data collection and the support options, see IBM Storage Insights documentation.

Also, refer to the following Redbooks publications:

- Intelligent Storage Management with Artificial Intelligence for IT Operations by using IBM Storage Insights, REDP-5755
- *Maximize Your Storage Investment: Use IBM Storage Insights for Observability and Artificial Intelligence for IT Operations*, SG24-8584

12.12.3 Case interaction with IBM Storage Insights

With IBM Storage Insights, customers can view and manage every case opened for any of their DS8000 storage systems. 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 Systems** → **Overview** page for the respective DS8000. On the right, under **System Actions**, select **Tickets** to view all open tickets and create a new one if needed. Alternatively, as shown in Figure 12-18, you can select **Get Support** to open a ticket directly.

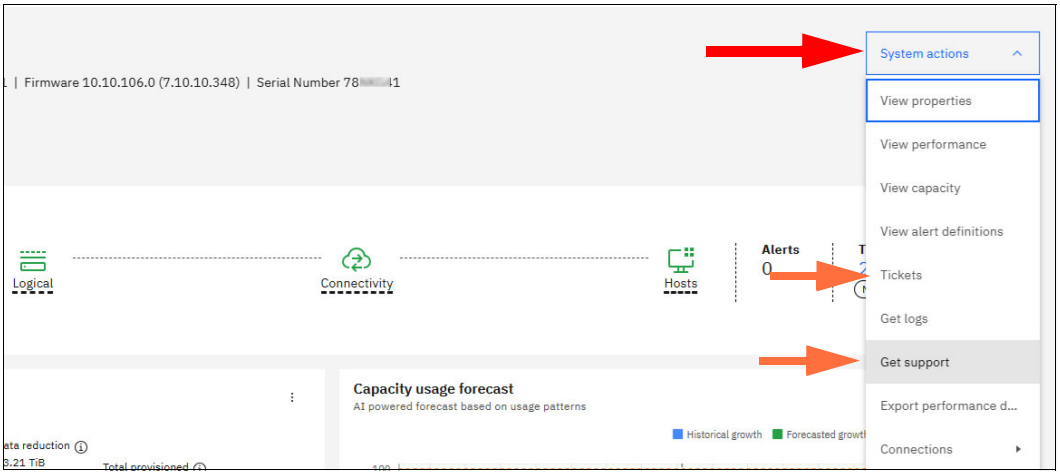


Figure 12-18 IBM Storage Insights: Storage system overview showing System actions

2. A new window opens with two options to choose from. Click **Create ticket**, as shown in Figure 12-19, then **Next**.

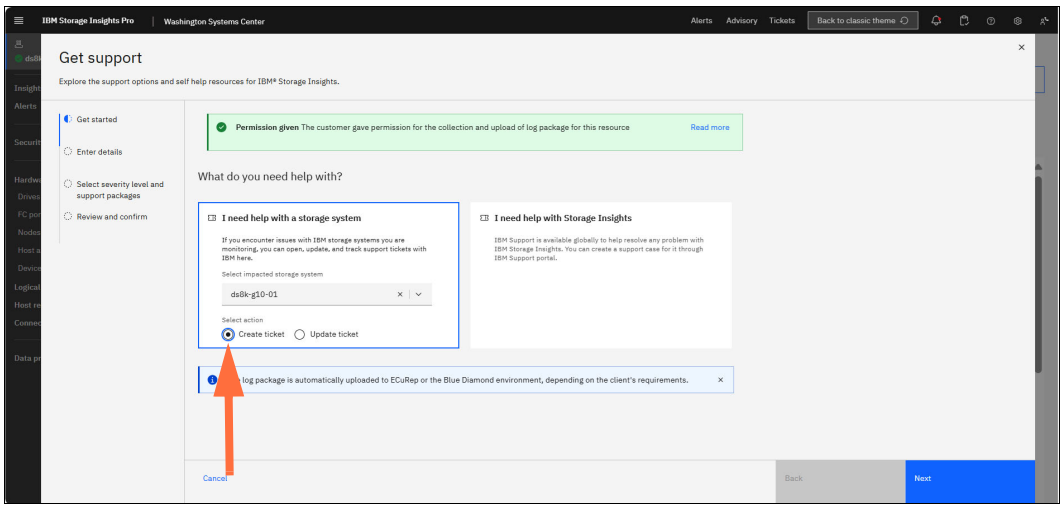


Figure 12-19 IBM Storage Insights: Create ticket

3. Figure 12-20 on page 456 shows the process of collecting the ticket information needed, which includes entering details about the problem. Click **Next**.

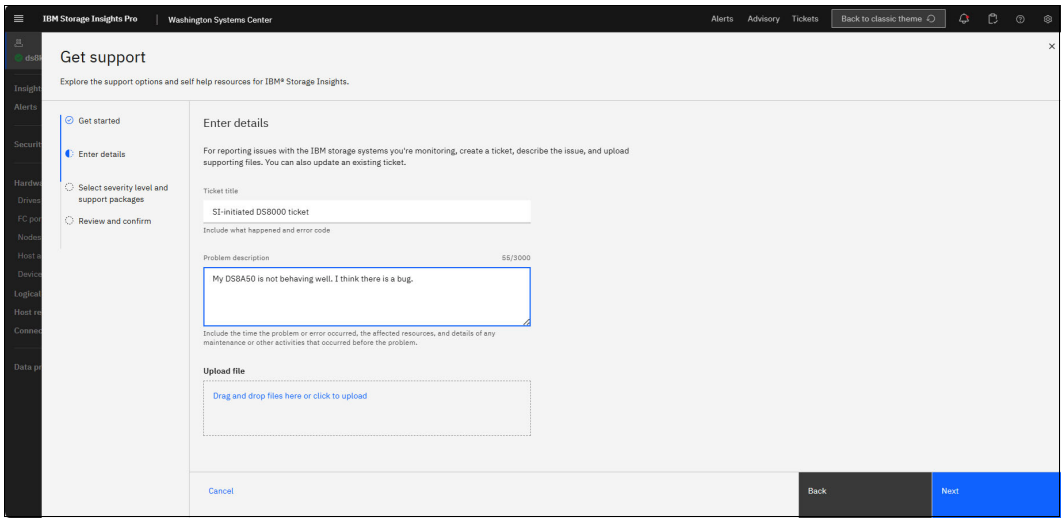


Figure 12-20 IBM Storage Insights: Entering ticket details

4. In the next window (Figure 12-21), you can select the appropriate severity category for your ticket based on its business impact. You also have the option to attach a support package. Creating a new PE package with the ticket is often helpful for accelerating resolution. If you are unsure which package to use, a general PE package is typically sufficient. However, if the issue involves the GUI or DSCLI, the ESSNI/DSGUI/API PE package can provide additional value.

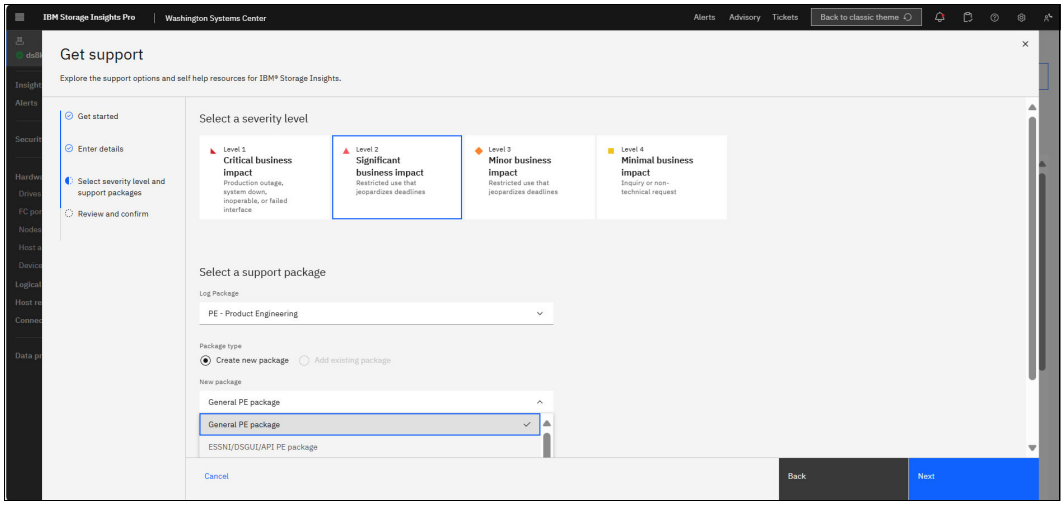


Figure 12-21 IBM Storage Insights: Selecting ticket severity and support package

5. Click **Next**. In the window shown in Figure 12-22 on page 457, you can review the ticket details once more and enter the correct contact information.

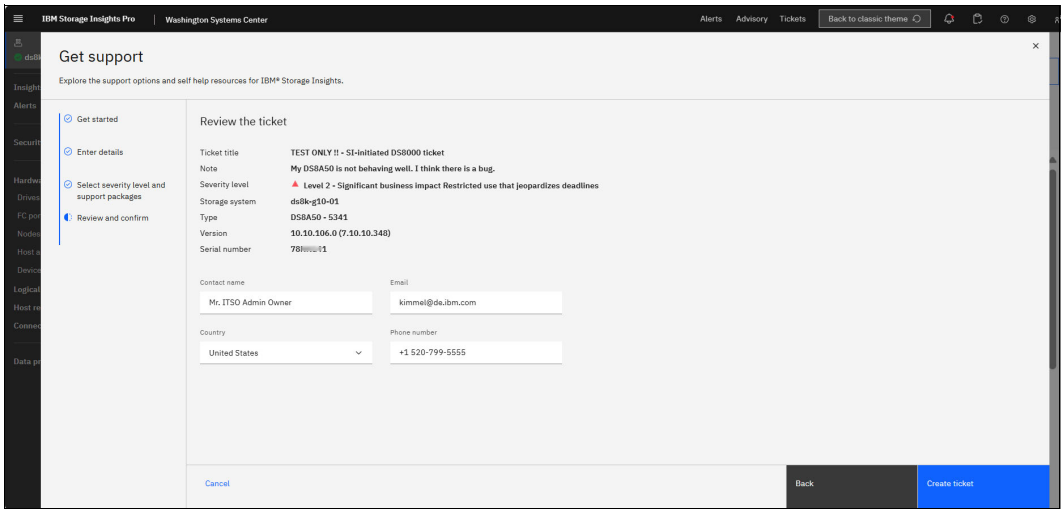


Figure 12-22 IBM Storage Insights, creating ticket: Final review

6. Click the blue **Create ticket** button to complete the process..

Updating IBM tickets

Whenever a customer needs to update a case or add additional information to an existing ticket, this can be done directly from IBM Storage Insights. As shown in Figure 12-19 on page 455, instead of selecting **Create ticket**, choose the **Update ticket** option. In the next window, you can either select from the list of existing open tickets or enter the ticket number you want to update.

12.12.4 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.

1. To set alert policies, in the left-hand side menu select **Administration** → **Alert Policy** when in the new theme; or using the classic theme, select **Configuration** → **Alert Policies**. Both options are shown in Figure 12-23 on page 458.

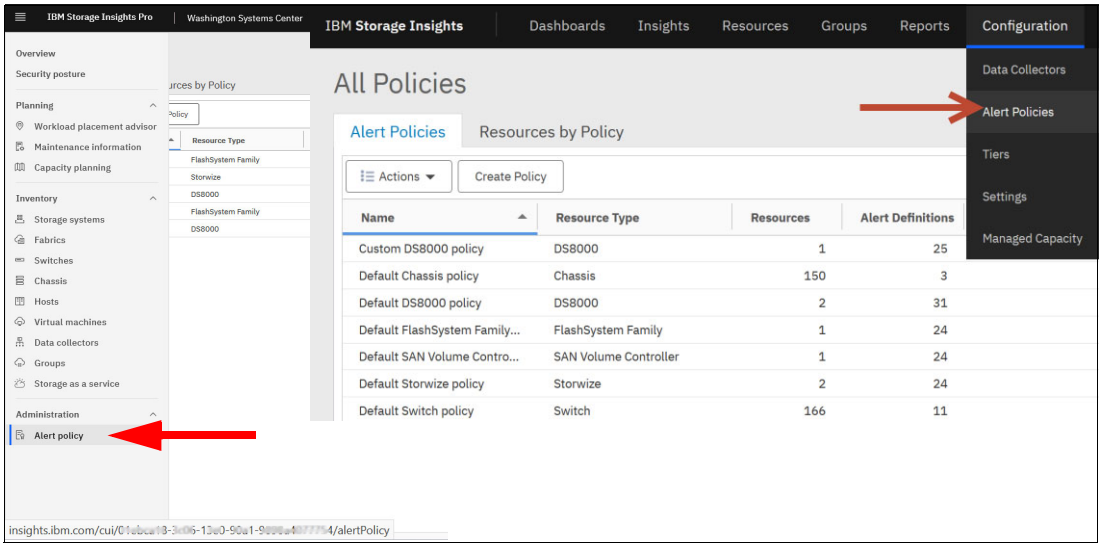


Figure 12-23 IBM Storage Insights: Alert Policies

2. There are many predefined default policies, but you can also 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 12-24.

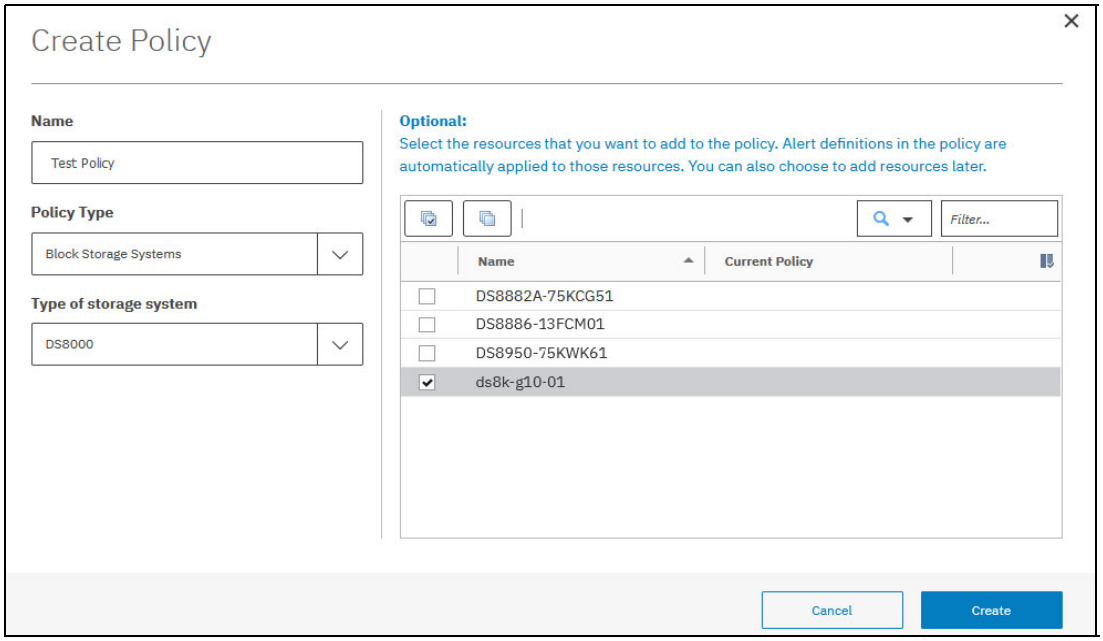


Figure 12-24 IBM Storage Insights: Create Policy

3. Select **Create**, and the policy is created. Then, define the alert definitions and save the changes, as shown in Figure 12-25.

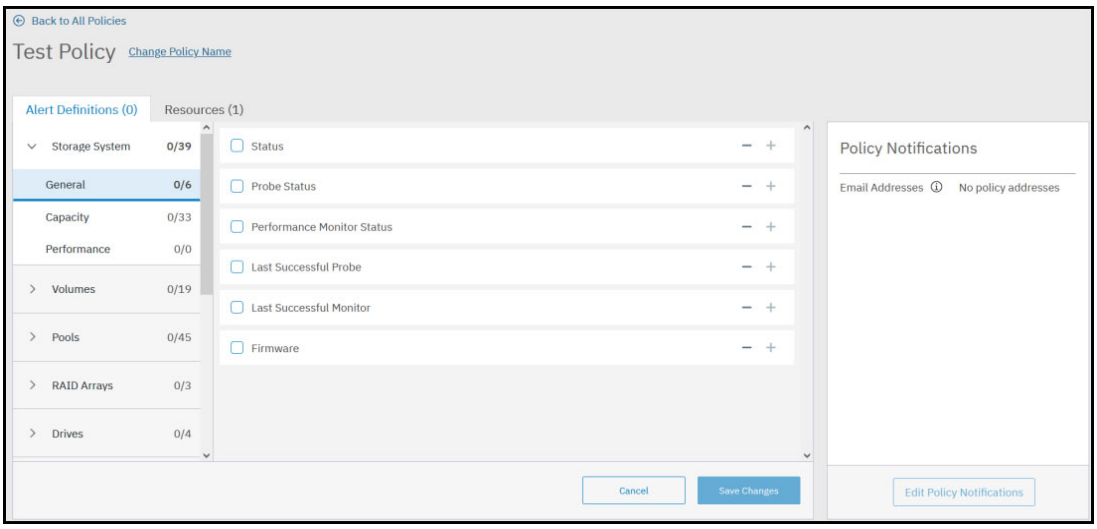


Figure 12-25 IBM Storage Insights: Alert Definitions

12.13 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`.

Figure 12-26 on page 460 shows the Call Home Connect Cloud web application. Alongside with your hosts and servers, all IBM devices can be listed on this page, allowing a quick view of their status. Code information, code recommendations, the warranty status, heartbeat status to IBM, machine-specific alerts, or the number of open support tickets is shown.

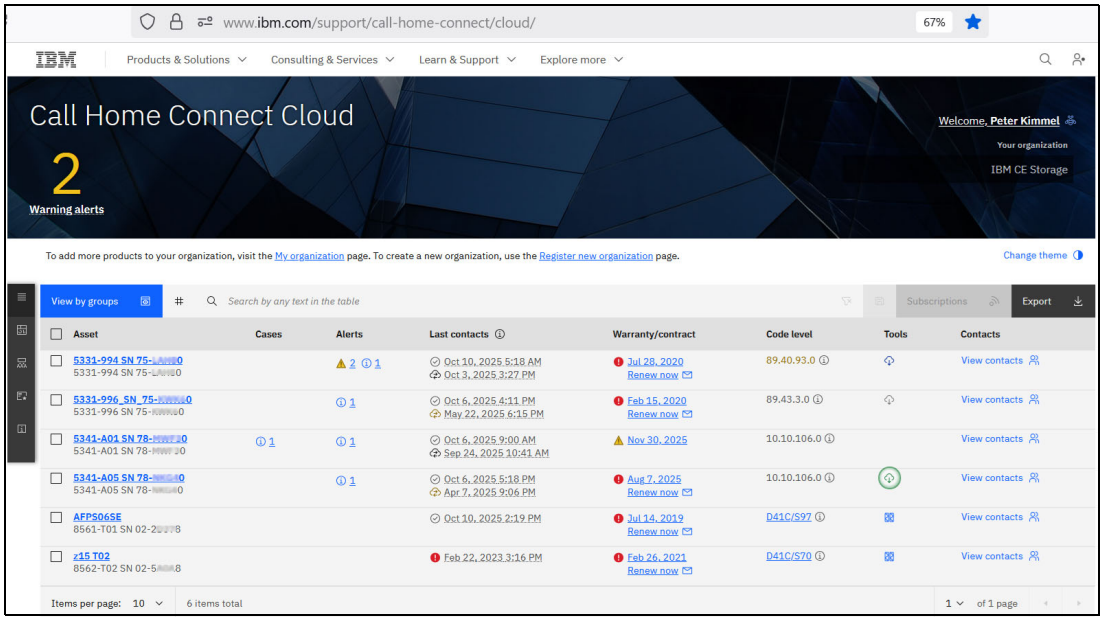


Figure 12-26 Call Home Connect Cloud web application

For the alerts or tickets listed, you can do a further drill-down. Figure 12-27 for instance shows the number and current status of the given open ticket, and from there you could move even further to see the full content of that one support case.

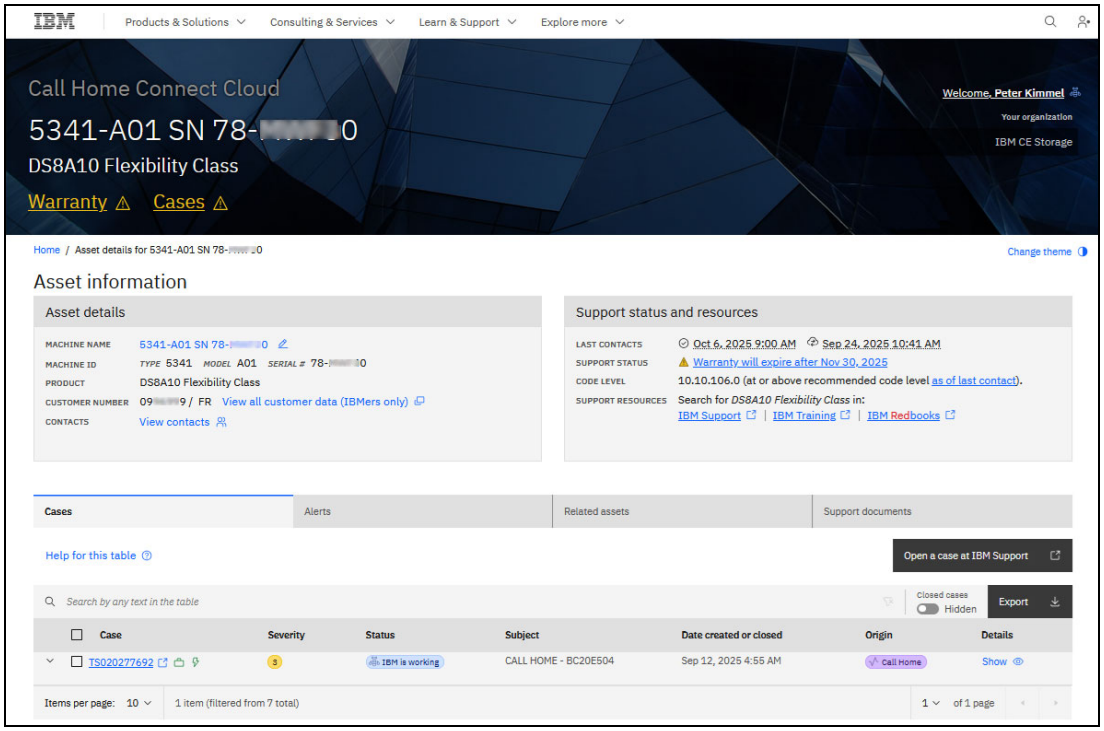


Figure 12-27 Call Home Connect Cloud: Information about open tickets

13

Optimizing performance: Practical tips and references

This chapter provides an overview of DS8A00 performance considerations and tips for configuration, host attachment, and planning for performance optimization.

We will discuss a wide range to topics related to the performance of the DS8A00 storage systems. This chapter has the following sections:

- “Data flow management and control” on page 462
- “Planning and analyzing arrays and RAID” on page 464
- “Understanding your workload” on page 467
- “Performance planning tool: IBM Storage Modeller” on page 473
- “Hardware planning for performance” on page 473
- “IBM Z host performance” on page 475
- “IBM i host performance” on page 478
- “Host adapters” on page 481
- “Practical performance management” on page 485

© Copyright IBM Corp. 2025.

13.1 Data flow management and control

The performance of a storage system is measured by how quickly it responds to read or write requests from an application server. This speed and accuracy depend on the design of the hardware, firmware, and system configuration.

Because data flows continuously from one component to another, minimizing bottlenecks is critical. Bottlenecks occur when data must wait for a busy component, reducing overall efficiency. Ideally, optimal use of system resources will reduce both the frequency and severity of these delays.

Key design considerations include:

- Balanced hardware components
- Appropriately sized storage system for the workload
- Proper logical configuration

Easy Tier® automatic rebalancing and tiering options can help achieve and maintain optimum performance even in an environment of ever-changing workload patterns, but they cannot replace a correct sizing of the storage system.

Although raw performance numbers can demonstrate that new hardware offers superior capabilities compared to previous generations or competing products, they often fail to reflect real-world production requirements. Maximum throughput figures are typically achieved under controlled test conditions by overloading the storage system. While these results may appear impressive, they overlook the trade-offs in response time introduced during such tests.

When planning a configuration to meet the requirements of a specific environment, consider the performance numbers for the individual flash drives, host adapters, and other components of the various models, and for the storage system as a whole. The IBM DS8000 family of storage servers uses the most current technology available and employs a rigorous testing approach to ensure a balanced, efficient, effective system.

13.1.1 Choosing hardware

Although general guidelines can help simplify hardware selection, they do not guarantee that the chosen configuration will meet the specific needs of your environment. These broad recommendations may speed up decision-making, but they can also lead to systems that fail to deliver the required performance. For example, a system might function adequately under normal conditions, yet struggle with throughput or latency during peak workloads or batch processing.

To avoid these scenarios, it is better to make an informed decision about your hardware requirements. IBM's performance modeling tool, Storage Modeller (StorM), is designed for this purpose. StorM calculates the latency (response time) and peak throughput (MB/sec or GB/sec) of the proposed storage server configurations based on a modeled workload. In addition, it estimates the expected internal utilization percentages, showing how loaded the system components will be under peak conditions. By using this tool, you can ensure that your configuration can handle your peak workload without needing to order excessive hardware based solely on generalized rules or recommendations.

See 13.4, "Performance planning tool: IBM Storage Modeller" on page 473 for more details.

Starting drive configurations

Once you have determined the overall capacity of the DS8A00 storage system, choosing the size of the FCMs is the next important decision. See Table 13-1 and Table 13-2 for starter configuration guidelines. Table 13-1 shows the DS8A00 single-frame starter configurations.

Table 13-1 Starter configuration examples for DS8A10 and single-frame DS8A50 or DS8A80

Configuration size	Industry standard		Configuration size	FlashCore Module
<36 TB	1.92 TB drives (1.6 TB for specific workloads)		<93 TB	4.8 TB drives
37 TB to 73 TB	1.92 TB or 3.84 TB drives		94 TB to 186 TB	4.8 TB or 9.6 TB drives
73 TB or 146 TB	3.84 TB or 7.68 TB drives		186 TB to 374 TB	9.6 or 19.2 TB drives
147 TB to 294 TB	7.68 or 19.2 TB drives (even if not exploiting compression)		374 TB to 1 PB	19.2 TB drives
294 TB to 1 PB	19.2 TB drives (even not exploiting compression)		>1 PB	Use DS8A50F unless no thin provisioning including FlashCopy Safeguarded Copy
> 1PB	Use DS8A50 unless no thin provisioning including FlashCopy Safeguarded Copy			

Table 13-2 shows the DS8A00 dual-frame starter configurations.

Table 13-2 Starter configuration examples for dual-frame DS8A50 or DS8A80

Configuration size	Industry standard		Configuration size	FlashCore Module
<73 TB	1.92 TB drives (1.6 TB for specific workloads)		<186 TB	4.8 TB drives
73 TB or 146 TB	1.92 TB or 3.84 TB drives		187 TB to 374 TB	4.8 TB or 9.6 TB drives
147 TB to 294 TB	3.84 TB or 7.68 TB drives		374 TB to 743 TB	9.6 TB or 19.2 TB drives
294 TB to 588 TB	7.68 TB or 19.2 TB drives (even if not exploiting compression)		>743 TB	19.2 TB drives
>588 TB	19.2 TB drives (even if not exploiting compression)			

13.1.2 Workload allocation to hardware

There are two approaches to allocating workloads to the hardware in a DS8A00. Each has its own benefits and drawbacks.

- Spreading workloads

  Spreading workloads involves multiple hosts using the same Host Adapter (HA) cards:

  - Pros: Reduced Total Cost of Ownership (TCO) and more effective use of hardware.
  - Cons: Potential for a single workload to dominate shared resources.

- Isolating workloads

  Isolating workloads involves dedicating specific hardware components to a single workload:

  - Pros: Components are available immediately when a workload demands service.
  - Cons: Certain components remain unused when their workload isn't demanding service.

For more information on hardware components and considerations for performance, see 13.5, "Hardware planning for performance" on page 473.

13.2 Planning and analyzing arrays and RAID

This section explains how to design and configure RAID arrays and extent pools to optimize performance, capacity utilization, and fault tolerance in DS8000 systems.

13.2.1 Logical configuration overview

Logical configuration refers to the creation of Redundant Array of Independent Disks (RAID) arrays and pools, and the assignment of the configured capacity to servers. The logical configuration can be created, applied, and modified by using the DS8000 Graphical User Interface (DS GUI) or DS8000 Command-line Interface (DS CLI). This is a customer responsibility.

Below is a brief overview of how virtualization builds upon the flash physical drives as a series of layers:

- Array site: Groups of eight identical physical flash drives, assigned automatically at install time.
- Array: RAID level, array type, and array configuration defined for a single array site
- Rank: Logical representation of the physical array that is formatted for use as FB or CKD storage types.
- Extents: Partitions of the capacity of an array the capacity of which depend upon the extent type, FB or CKD.
- Extent pools: Logical construct to aggregate the extents from a set of ranks, and it forms a domain for extent allocation to a logical volume.

For more detailed information on logical configuration, see *IBM DS8900F Performance Best Practices and Monitoring*, SG24-8501.

13.2.2 Logical configuration performance considerations

This section outlines key factors and best practices for configuring RAID arrays and ranks to ensure balanced performance, minimize bottlenecks, and meet workload-specific requirements in DS8000 systems.

Arrays and RAID configurations

Array sites are dynamically created and assigned to HPFEs and their DA pairs during DS8000 installation. The DS Storage Manager GUI automatically distributes arrays across all DAs to achieve balanced performance, and initial formatting through the GUI typically delivers optimal results for most environments. For specific requirements, such as isolating arrays by DA pairs, the DSCLI can be advantageous because it offers additional configuration options. Arrays and ranks can be planned together since they have a one-to-one relationship, but they must be created separately.

Note: Starting with DS8A00 RAID 5 is no longer supported. Valid RAID configurations are RAID 6 and RAID 10.

RAID 10 optimizes high performance while maintaining fault tolerance for disk drive failures. The data is striped across several disks, and the first set of disk drives is mirrored to an identical set. RAID 10 can tolerate at least one, and in most cases, multiple disk failures if the primary copy and the secondary copy of a mirrored disk pair do not fail at the same time. Still, RAID 6 offers more usable capacity than RAID 10 by providing an efficient method of data protection in double disk errors, such as two drive failures, two coincident medium errors, or a drive failure and a medium error during a rebuild.

For random-write I/O operations, RAID 10 and RAID 6 differ significantly in performance characteristics. With RAID 10, each write operation at the disk back end triggers two I/O operations to the rank. In contrast, RAID 6, with two parity blocks, incurs a higher write penalty, requiring up to six back-end I/O operations for a single random small-block write. Although this represents a worst-case scenario, it provides a useful estimate of the expected back-end workload and helps in planning the appropriate number of arrays.

On modern disk systems such as the DS8000, write operations are cached by the storage subsystem and handled asynchronously, resulting in short write response times for attached host systems. This means any RAID 6 write penalties are effectively masked from the host in terms of disk response time. Typically, a write request sent to the DS8000 is first written into the storage server cache and persistent cache, and the I/O operation is immediately acknowledged to the host as complete. As long as there is available space in these cache areas, the application only experiences the time required to place data into the cache. Therefore, whether RAID 6 or RAID 10 is used has no impact on the perceived response time.

Occasionally, the cache can fill up with no space for the next write request. When this occurs, the storage server signals the host system to retry the I/O write operation. In the time that it takes the host system to retry the I/O write operation, the storage server likely can destage part of the data, which provides free space in the cache and allows the I/O operation to complete on the retry attempt.

High-Capacity versus High-Performance ranks

The performance of High-Capacity and High-Performance ranks is generally comparable. However, it is important to note that High-Capacity ranks offer lower internal parallelism compared to their High-Performance counterparts.

Large versus small extents

Extent sizes depend on the extent type, either FB or CKD. You can decide whether to use large or small extents.

- FB large extent: 1 Gib (2^{30} bytes)
- FB small extent: 16 MiB (16 x 2^{20} bytes)
- CKD large extent: 1113 cylinders (1 3390 Model 1)
- CKD small extent: 21 cylinders

Whether you use small or large extents depends on the goals you want to achieve. Small extents provide better capacity utilization, particularly for thin-provisioned storage, so they are generally preferred for thin-provisioned volumes. However, managing many small extents can cause minor performance degradation during initial allocation. For example, formatting 1 GB requires one storage allocation with large extents but 64 allocations with small extents. That said, small extents often offer some operational performance advantages due to greater distribution, increased parallelism, and potentially more ranks involved, which can improve throughput in many cases. Therefore, IBM typically recommends small extents as the default —especially if you plan to use Safeguarded Copy -unless you are working with extremely large multi-terabyte assignments.

Release extents

Although releasing extents is part of a Copy Services activity, it can temporarily consume back-end resources. To minimize performance impact, IBM recommends performing this task during periods of low system usage, especially when other resource-intensive operations are scheduled.

Extent pools

Single-tier or homogeneous extent pools contain ranks from only one storage class or tier. These pools can be either single-rank or multi-rank.

Single-rank homogeneous extent pools introduce limitations, such as reduced flexibility and potential performance constraints. In today's complex, mixed environments, single-rank pools are generally not recommended unless strict isolation and fine-grained performance control are required for a specific workload.

If not managed properly, these pools can lead to workload skew and rank hot spots, which negatively impact overall system performance.

With a homogeneous multi-rank extent pool, you take advantage of the advanced DS8000 virtualization features to spread the workload evenly across the ranks in an extent pool to achieve a well-balanced data distribution with considerably less management effort. Multi-rank extent pools that use storage pool striping are the general configuration approach today on modern DS8A00 storage systems to spread the data evenly across the ranks in a homogeneous multi-rank extent pool and thus reduce skew and the likelihood of single-rank hot spots.

Multi-tier or hybrid extent pools

Multi-tier or hybrid extent pools consist of ranks from different storage classes or tiers. Easy Tier can automatically manage data placement within and across these tiers, optimizing both performance and cost. However, since the DS8A00 series, multi-tier pools have become extremely rare in real-world deployments, with most configurations using single-tier pools. Internal tiering still occurs within each FlashCore Module (FCM), which combines fast SLC flash with more cost-efficient QLC flash inside the same drive.

Multi-tier extent pools are especially suited for mixed, resource-sharing workloads. Tiered storage is an approach that uses multiple types of storage across the infrastructure—a mix of higher-performing/higher-cost storage and lower-performing/lower-cost storage—placing data based on its specific I/O access characteristics. When correctly balanced, tiered storage delivers the lowest cost and best performance solution.

Best practices summary:

- Use RAID 10 for workloads requiring high performance and fault tolerance; choose RAID 6 for better capacity efficiency with strong protection against double disk failures.
- Prefer multi-rank extent pools to reduce workload skew and avoid rank hot spots.
- Use small extents for thin-provisioned volumes and Safeguarded Copy. Large extents are generally chosen for faster initial allocation, but the main driver for using large extents is extremely large capacity requirements—such as configurations needing more than six drive sets of 19.2 TB FCMs. For most clients, initial allocation speed is not a major concern.
- Perform extent release during low-usage periods to minimize performance impact.
- Leverage Easy Tier for automated tiering in hybrid pools to optimize performance and cost.

13.3 Understanding your workload

This section describes common workload types and their I/O characteristics, helping you accurately profile, isolate, and optimize workloads for maximum performance on DS8000 systems.

13.3.1 General workload types

This section provides an overview of common workload patterns and their I/O characteristics, serving as a foundation for performance planning and optimization in DS8000 environments.

Typical online transaction processing workload

This workload is characterized by mostly the random access of small-sized I/O records (less than or equal to 16 KB) with a mix of 70% reads and 30% writes. This type of workload is typical because it is the basis for most of the benchmark and performance tests.

Microsoft Exchange Server workload

This type of workload can be similar to an OLTP workload, but it has a different read/write balance. It is characterized by many write operations, up to 60%, with high random numbers. Also, the size of the I/O can be high, with blocks up to 128 KB, which is explained by the nature of the application.

Sequential workload

Sequential access is one of the original workload types for data storage. Sequential workload is good for prefetch and putting data into cache because blocks are accessed one by one, and the storage system can read many blocks to unload the flash modules or disks. Similarly, write requests work well because the storage system can optimize the access to the flash modules disks and write several tracks or cylinders at a time.

Batch jobs workload

Batch workloads have several common characteristics:

- Mixture of random database access, skip-sequential, pure sequential, and sorting.
- Large block sizes up to 128 - 256 KB.
- High volume of write activity, and read activity.
- The same volume extents might scramble for write and read at the same time.
- Batch workloads include large data transfers and high path utilizations.
- Batch workloads are often constrained to operate within a particular window of time when online operation is restricted or shut down. Poor or improved performance is often not recognized unless it affects this window.

Plan when to run batch jobs: Plan all batch workload activity for the end of the day or at a slower time of day. Normal activity can be negatively affected with the batch activity

Sort jobs workload

Most sorting applications are characterized by large transfers for input, output, and work data sets. DFSORT is IBM's high-performance sort, merge, copy, analysis and reporting product and is an optional feature of z/OS.

Read-intensive cache friendly and unfriendly workloads

Use the cache hit ratio to estimate the cache-friendly read workload. If the ratio is more than 50%, the workload is cache friendly. If you have two serviced I/Os to the same data and one I/O is serviced from cache, it is a 50% cache hit ratio. Write workload, and sequential workload are always cache friendly.

Cache unfriendly workloads include batch processing and reporting, as well as high random-access numbers. Though it may not always be easy to divide known workload types into cache friendly and cache unfriendly designations, the DS8000 storage systems use the following powerful read-caching algorithms to deal with cache unfriendly workloads:

- Sequential Prefetching in Adaptive Replacement Cache (SARC)
- Adaptive Multi-stream Prefetching (AMP)
- Intelligent Write Caching (IWC)

Table 13-3 provides a summary of the characteristics of the various types of workloads.

Table 13-3 Workload types

Workload type	Characteristics	Representative of this type of process
Sequential read	Sequential 128 - 1024 KB blocks Read hit ratio: 100%	Database backups Batch processing Media streaming
Sequential write	Sequential 128 - 1024 KB blocks Write ratio: 100%	Database restores and loads Batch processing File access
z/OS cache uniform	Random 4 KB record R/W ratio 3.4 Read hit ratio 84%	Average database CICS/VSAM IBM IMS

Workload type	Characteristics	Representative of this type of process
z/OS cache standard	Random 4 KB record R/W ratio 3.0 Read hit ratio 78%	Representative of typical database conditions
z/OS cache friendly	Random 4 KB record R/W ratio 5.0 Read hit ratio 92%	Interactive Existing software
z/OS cache hostile	Random 4 KB record R/W ratio 2.0 Read hit ratio 40%	Db2 logging
Open read-intensive	Random 4 - 8 - 16 - 32 KB record Read% = 70 - 90% Hit ratio 28%, 1%	Databases (Oracle and Db2) Large DB inquiry Decision support Warehousing
Open standard	Random 4 - 8 - 16 - 32 KB record Read% = 70% Hit ratio 50%	OLTP General Parallel File System (IBM GPFS, now known as IBM Storage Scale)

13.3.2 Database workload

Database workload depends on the application that is written for the database and the type of work that this application performs. The database environment is often difficult to typify because I/O characteristics differ greatly. Users must understand the unique characteristics of their database capabilities before generalizing the performance.

Query workload

- High read content
- Mix of write and read content
- Random access, and sequential access

A well-tuned database keeps characteristics of the queries closer to a sequential read workload.

Logging workload

There are several types of log files in the database:

- Online transaction logs are used to restore the last condition of the database when the latest transaction failed.
- Archive transaction logs are used to restore a database state up to the specified date.

The workload pattern for logging consists mostly of sequential writes. Database performance heavily depends on writes to the online transaction logs. If the database is highly write-intensive, consider using RAID 10. Additionally, if the configuration includes several extent pool pairs in its logical layout, consider physically separating the log files from the flash modules or disks that store the data and index files.

Transaction environment workload

Database transaction workloads have these characteristics:

- Low to moderate read hits, depending on the size of the database buffers
- Cache unfriendly for certain applications

- Deferred writes that cause low write-hit ratios, which means that cached write data is rarely required for reading
- Deferred write chains with multiple locate-record commands in chain
- Low read/write ratio because of reads that are satisfied in a large database buffer pool
- High random-read access values, which are cache unfriendly

A database can benefit from using a large amount of server memory for the large buffer pool. For example, the database large buffer pool, when managed correctly, can avoid a large percentage of the accesses to flash or disk. Depending on the application and the size of the buffer pool, this large buffer pool can convert poor cache hit ratios into synchronous reads in Db2. You can spread data across several RAID arrays to increase the throughput even if all accesses are read misses. Db2 administrators often require that table spaces and their indexes are placed on separate volumes. This configuration improves both availability and performance.

Utilities workload

Database utilities, such as loads, reorganizations, copies, and recovers, generate high read and write sequential and sometimes random operations. This type of workload takes advantage of the sequential bandwidth performance of the back-end storage connection, such as the PCI-Express bus for the DA pairs, and the use of flash drives and Easy Tier automatic mode enabled.

13.3.3 Application workload

There are four typical categories of applications, according to their I/O behavior:

- Need for high throughput
- Need for high throughput and a mix of read/write (R/W)
- Need for high I/O rate and throughput
- Need for high I/O rate

Table 13-4 summarizes these workload categories and common applications.

Table 13-4 Application workload types

Category	Application	Read/write ratio	I/O size	Access type
4	General file serving	Expect 50/50	64 - 256 KB	Sequential mostly because of good file system caching.
4	Online transaction processing	50/50, 70/30	4 KB, 8 KB, or 16 KB	Random mostly for writes and reads. Bad cache hits
4	Batch update	Expect 50/50	16 KB, 32 KB, 64 KB, or 128 KB	Almost 50/50 mix of sequential and random. Moderate cache hits.
1	Data mining	90/10	32 KB, 64 KB, or larger	Mainly sequential, some random.

Category	Application	Read/write ratio	I/O size	Access type
1	Video on demand	100/0	256 KB and larger	Sequential, good caching.
2	Data warehousing	90/10, 70/30, or 50/50	64 KB or larger	Mainly sequential, rarely random, and good caching.
2	Engineering and scientific	100/0, 0/100, 70/30, or 50/50	64 KB or larger	Sequential mostly, good caching.
3	Digital video editing	100/0, 0/100, or 50/50	128 KB, 256 - 1024 KB	Sequential, good caching.
3	Image processing	100/0, 0/100, or 50/50	64 KB, 128 KB	Sequential, good caching.
1	Backup, restore	100/0, 0/100	256 - 1024 KB	Sequential, good caching.

13.3.4 Profiling workloads in the design phase

Assessing the I/O profile before you build and deploy the application requires methods of evaluating the workload profile without measurement data. In these cases, as a preferred practice, use a combination of general rules based on application type and the development of an application I/O profile by the application architect or performance architect. This approach helps you make informed decisions about storage configuration, RAID level, and capacity planning early in the design phase, reducing the risk of performance bottlenecks after deployment.

For general rules for application types, see Table 13-3 on page 468.

The following requirements must be considered when developing an I/O profile for an application:

- User population: Total users, average active users, and peak active users.
- Transaction distribution: Iterations per user, I/O per transaction, and I/O type.
- Logical I/O profile: Derived from user population and transaction distribution.
- Physical I/O profile: Detailed breakdown of actual I/O operations, including size, frequency, and access patterns.

For a more detailed breakdown of these assessments, see *IBM DS8900F Performance Best Practices and Monitoring*, SG24-8501.

13.3.5 Workload isolation versus resource sharing

This section compares strategies for dedicating hardware resources to specific workloads versus sharing resources across multiple workloads, highlighting their impact on performance, availability, and flexibility in DS8000 environments.

Workload isolation

With workload isolation, a high-priority workload uses dedicated DS8000 hardware resources to reduce the impact of less important workloads. Workload isolation can also mean limiting a

lower-priority workload to a subset of the DS8000 hardware resources so that it does not affect more important workloads by fully using all hardware resources.

Isolation provides ensured availability of the hardware resources that are dedicated to the isolated workload. It removes contention with other applications for those resources. Typically, isolation is implemented to improve the performance of certain workloads by separating different workload types.

Flash drive capacity isolation

This section explains how isolating flash drive capacity at different levels, such as rank, extent pool, or device adapter, can help ensure predictable performance for critical workloads in DS8000 systems.

At the rank level, volumes for a single workload are allocated to specific ranks, which may be separated by disk type, RAID type, or storage type (FB or CKD).

At the extent pool level, isolation can be implemented by assigning ranks to dedicated pools for a workload.

At the Device Adapter (DA) or High-Performance Flash Enclosure (HPFE) level, all ranks on one or more DA pairs or an HPFE pair can be dedicated to a workload. This approach is recommended only for extremely large workloads, those requiring multiple gigabytes per second of throughput and involving heavy, large-block sequential activity.

At the processor complex level, all ranks assigned to extent pools managed by a specific processor complex (0 or 1) can be dedicated to a workload. However, this method is generally not preferred because it effectively halves the available resources and bandwidth for the workload.

At the storage unit level, all ranks within a physical DS8A00 storage system can be dedicated to a workload. Additionally, pinning volumes to a tier in hybrid pools allows for a simpler extent pool structure while prioritizing volumes by assigning them to specific tiers. These pinned volumes can be dynamically changed at any time.

Host connection isolation

This section explains how dedicating specific I/O ports, host adapters, or enclosures to individual workloads can reduce resource contention and improve performance in DS8000 environments.

At the I/O port level, certain DS8A00 ports can be dedicated to a specific workload. This is a common practice because ports inherently isolate FCP and FICON traffic. It is also typically used to separate remote mirroring traffic (FCP) from Open Systems host traffic (FCP).

At the Host Adapter (HA) level, dedicating specific HAs to a workload is considered a preferred practice, especially for separating FCP and FICON traffic across different HBAs, although this is not strictly required. Workloads with heavy, large-block sequential activity should be prioritized for HA isolation to ensure optimal performance.

At the I/O enclosure level, isolation is technically possible but generally unnecessary, as the benefits rarely justify the added complexity.

Workload resource sharing

Multiple resource-sharing workloads can have logical volumes on the same ranks and can access the same DS8A00 HAs or I/O ports. Resource-sharing allows a workload to access more DS8A00 hardware than can be dedicated to the workload, providing greater potential

performance, but this hardware sharing can result in resource contention between applications that impacts overall performance at times.

Easy Tier extent pools typically are shared by multiple workloads because Easy Tier with its automatic data relocation and performance optimization across multiple storage tiers provides the most benefit for mixed workloads.

13.4 Performance planning tool: IBM Storage Modeller

IBM Storage Modeler (StorM) is a sizing and modeling tool for IBM Storage Systems, based on an open client/server architecture, prepared to be connected to other tools, using a web-based graphical user interface. Its ability to collaborate by sharing project files is a huge differentiator, and the interactive canvas enables easy design and visualization of a solution architecture. IBM Storage Modeler is platform agnostic and runs on most current web browsers.

13.5 Hardware planning for performance

This section provides an overview of the DS8A00 hardware components and their impact on system performance.

13.5.1 Host Adapters

Host Adapters (HA) operate on the “front-end” of the DS8000 and serve two functions:

- They provide I/O connectivity between the host server and the DS8000 storage system for reads and writes.
- They carry synchronous (Metro Mirror) and asynchronous (Global Mirror and Global Copy) replication between source and target DS8000s.

> **Note:** Although host I/O can be either FCP or FICON, replication between DS8000s is exclusively via FCP.

13.5.2 Device Adapters

Device Adapters (DAs), also known as Virtual RAID Adapters (VRA), also serve two functions, but on the back-end of the DS8000.

1. They manage RAID protection and sparing for the flash arrays or rank.
2. They manage read and write I/Os to and from those ranks, respectively.

> **Note:** Although VRAs are now ACTIVE/ACTIVE it remains a best practice to split resources in a way that maintains physical redundancy throughout the hardware.

13.5.3 High-Performance Flash Enclosure - Gen3

A DS8000 High-Performance Flash Enclosure - Gen3 (HPFE) is a 2U storage enclosure that houses a redundant pair of ESMs and one, two, or three flash drivesets. HPFEs are installed

in pairs. A single flash driveset consists of 16 industry-standard flash drives or IBM FlashCore Modules, with eight drives installed in each enclosure.

13.5.4 zHyperLink

zHyperLink is a point-to-point optical connection designed to provide ultra-low response times (latency) between IBM Z servers and DS8000 storage systems. zHyperLink does not replace FICON connectivity. Rather, it reduces I/O latency, for selected eligible I/Os, to the point that z/OS no longer needs to undispatch the running task.

zHyperLinks are typically added in pairs for redundancy and have a point-to-point distance limitation of 150 m (492 feet). They require specific zHyperlink adapters on System Z, zHyperLink cabling, and zHyperLink adapters on the DS8A00. As with other connection types, we recommend spreading zHyperLinks across the frames as evenly as possible.

For environments requiring longer distances, such as PPRC (Peer-to-Peer Remote Copy) configurations, zHyperLink cannot be used due to its 150 m limitation. In these cases, IBM provides options like Consistent Read from Secondary, which enables applications to read from the secondary copy while maintaining data consistency.

13.5.5 Multipathing

Open Systems environments require multipathing software to fully leverage multiple paths for both high availability and load balancing. Each host should have at least two independent connections to separate Host Adapters (HAs) located in different I/O enclosures to eliminate single points of failure. Ensure that you provision enough HA cards to meet performance and throughput requirements. Additionally, deploying multiple switches or SAN directors in the fabric design further enhances resiliency and prevents a single component failure from impacting connectivity.

In the IBM Z environment, it is typical to have four paths installed from each host to the storage system. HCD or IOCDS define the LCU(s) accessible to the host system via these paths, while built-in utilities such as DPS and DPR provide multipath functionality and are built in to the IBM Z architecture.

13.5.6 Host attachment spreading

Spreading connections to HA ports should be approached depending on the characteristics of the workloads involved. For varied workloads, it may be better to isolate workload types to pairs of HAs. If similar hosts are running similar workloads, spreading connections out among the HA cards is a good approach. Keep the following general guidelines in mind when planning host connections:

- Spread the paths from all host systems across the available I/O ports, HA cards, and I/O enclosures, when possible.
- Spread the host paths that access the same set of volumes as evenly as possible across the available HA cards and I/O enclosures.
- Ensure that each host system uses a multipathing device driver and a minimum of two host connections to different HA cards in separate I/O enclosures on the DS8000 storage system. In R10.x systems, I/O bays are arranged in vertical pairs, with each bay occupying the full width of the rack. Within each pair, the even-numbered bay is positioned above the odd-numbered bay. For optimal availability and balanced workload distribution, host connections should be spread across multiple I/O enclosures and HA cards.

- Do not use the same DS8000 I/O port for host attachment and Copy Services remote replication (such as Metro Mirror or Global Mirror). Ideally, even separate off HA pairs (with all their ports) for Copy Services traffic only if enough adapters are available.
- Consider using separate HA cards for the FICON protocol and FCP. Even though I/O ports on the same HA can be configured independently for the FCP and the FICON protocol, it might be preferable to isolate your z/OS environment (FICON) from your Open Systems environment (FCP).

13.5.7 Caching

A number of factors must be considered to accurately determine the needed size of cache. However, it is a common approach to base the amount of cache on the amount of disk capacity, as shown in these common general rules:

- For Open Systems, each TB of drive capacity needs 1 GB - 2 GB of cache.
- For Z Systems, each TB of drive capacity needs 2 GB - 4 GB of cache.

13.6 IBM Z host performance

This section focuses on performance considerations for IBM Z environments connected to DS8000 systems. It covers monitoring tools, workload characteristics, and configuration best practices to ensure optimal throughput, low latency, and balanced resource utilization. You'll also find guidance on leveraging advanced technologies such as zHyperLink and Easy Tier for performance optimization.

13.6.1 Performance monitoring with RMF

The Resource Management Facility (RMF), a component of the z/OS operating system, provides detailed performance metrics for DS8000 storage systems as well as other storage resources. It provides performance information for the DS8000 and other storage systems. RMF can help with monitoring the following performance components:

- I/O response time
- IOP / SAP
- FICON host channel
- FICON director
- zHyperLink
- Symmetric multiprocessing (SMP)
- Cache and nonvolatile storage (NVS)
- Enterprise Storage Servers
  - FICON/Fibre port and host adapter (HA)
  - Extent pool and rank/array

RMF uses a long-term data collector for all types of resources and workloads, a snapshot data collector for address space states and resource usage, and a Short-term data collector for problem determination, workflow delay monitoring, and goal attainment supervision. Collected data is then stored and processed. It may be stored in SMF records, or in the case of short-term data (Monitor III), VSAM data sets. The RMF postprocessor can also extract historical data from the long-term data collector (Monitor I).

The RMF Postprocessor provides a number of reports that are beneficial to analyzing the performance of the DS8A00 storage server:

Direct Access Activity report

The RMF Direct Access Device Activity report is a good starting point in analyzing the storage systems' performance. It lists the activity of individual devices, and their average response time for the reporting period.

The device activity report provides various components of the overall response time for an I/O operation:

- IOS Queuing (IOSQ) time: Queuing at the host level.
- Pending (PEND) time: Response time of the storage hardware (fabric and HA).
- Disconnect (DISC) time: Stage or de-stage data to or from cache, and remote replication impact.
- Connect (CONN) time: Data transfer time.
- I/O Interrupt delay (INT): Time between finishing an I/O operation and interrupt.

I/O Queuing Activity report

The I/O Queuing Activity report will show how busy the I/O processors (IOPs) in a Z Systems server are. It will also show the Alias Management Group (AMG) details. An AMG is a set of logical subsystems (LSS) which are also known as Logical Control Units (LCU) or just control units. Finally, the report will break down I/O Queuing Activity by LCU.

FICON Host Channel report

FICON Channel statistics can be viewed in the FICON Host Channel report. This report gives insight into the utilization of the channel hardware related to the I/O activity on this LPAR (PART), the total utilization from all LPARs sharing the channel, and the usage of the internal bus that connects the channel to the CPC. Generally speaking, the total utilization should not exceed 50%, while the maximum BUS utilization should be around 40%.

FICON Director Activity report

The FICON director is the switch that connects the host FICON channel to the DS8000 HA port. The measurements provided in the report represent the total I/O through the individual ports on the director. The data are not broken down by LPAR or device.

The important performance metric is AVG FRAME PACING. This metric shows the average time (in microseconds) that a frame waited before it was transmitted. The higher the contention on the director port, the higher the average frame pacing metric. High frame pacing negatively influences the CONNECT time.

Cache and NVS report

The CACHE SUBSYSTEM ACTIVITY report is provided at the LCU level and provides useful information for analyzing the reason for high DISC time. The report breaks down the LCU level I/O requests by read and by write. It shows the rate, the hit rate, and the hit ratio of the read and the write activities.

Enterprise Disk Systems report

The Enterprise Disk Systems (ESS) report provides measurements of rank activity, FICON/FCP port and HA activity, as well as extent pool information.

13.6.2 Z Systems planning and configuration

An IBM Z system's workload is too complex to be accurately represented by a few numbers or general guidelines. Likewise, the performance of a storage system is more than the sum of its individual components. Advanced technologies such as zHyperLink and zHyperWrite can

significantly enhance overall throughput and impact additional functions like point-in-time copies or remote replication, which introduce extra workload considerations.

To account for these factors, IBM recommends modeling the new solution using StorM. For IBM Z environments, this modeling leverages RMF data and real workload characteristics, enabling precise comparisons between current and proposed configurations. StorM also allows you to factor in future growth (capacity and workload) and evaluate the impact of Copy Services and other IBM Z synergy features on performance.

13.6.3 Optimizing performance

Easy Tier provides several capabilities that you can use to solve performance issues. Here are some examples:

- Distribute workload across all resources within a storage tier
- Manage skewed workload in multitier (hybrid) extent pools
- Give new workloads a good start
- Add or remove resources

Easy Tier is designed to provide a balanced and stable workload distribution. It aims at minimizing the amount of data that must be moved as part of the optimization.

FICON channels

When planning your FICON channels, consider the following constraints and considerations:

- IBM Z systems support a maximum of eight channel paths to each logical control unit (LCU).
- A FICON channel port can be shared between several z/OS images and can access several DS8000 HA ports.
- A DS8000 Host Adapter (HA) port can be shared by multiple z/OS images and even by several IBM Z servers, enabling flexible connectivity and efficient resource utilization.
- A FICON host channel or DS8000 HA port has specific throughput capabilities that are not necessarily equivalent to the link speed they support.
- A Z Systems FICON feature or a DS8000 Host Adapter (HA) card has an overall throughput limit that may be lower than the combined theoretical bandwidth of all its individual ports. This means that even if multiple ports are active, the card's aggregate performance is constrained by its maximum design capacity.

.

Summary of best practices for IBM Z host performance:

- Leverage RMF reports (Direct Access, I/O Queuing, FICON Channel, Cache/NVS) to analyze latency, throughput, and resource utilization.
- Keep FICON Express CHPID utilization below 50% and FICON Bus busy below 40% for optimal performance.
- Monitor FICON Director activity. Watch AVG FRAME PACING metrics to detect contention and prevent increased CONNECT times.
- Size cache appropriately. For IBM Z workloads, allocate 2 - 4 GB of cache per TB of storage to maintain high hit ratios.
- Do not mix different FICON link speeds (8/16/32 Gbps) to the same LCU.
- Spread host paths across HA cards and enclosures. Distribute connections evenly for redundancy and balanced workload.
- Enable Easy Tier to manage skewed workloads and optimize performance across tiers.
- Deploy zHyperLink for critical workloads requiring sub-millisecond response times.
- Use IBM Storage Modeller (StorM) to validate performance under real workload conditions.

13.7 IBM i host performance

This section provides guidance on optimizing DS8000 performance in IBM i environments. It covers storage architecture, connectivity options, multipathing, and configuration best practices. You'll also find recommendations for RAID selection, LUN sizing, and extent pool design to ensure high availability and efficient I/O operations.

13.7.1 IBM i storage architecture

To understand the performance of the DS8000 storage system with the IBM i system, you need insight into IBM i storage architecture.

- Single level storage: Main memory and the flash unit space are treated as one unit
- Object-based architecture: Objects (files, programs, libraries, etc.) are packaged with usage rules and placed in a location that ensures the best performance. Normally, the object's data is spread across multiple flash units
- Storage management: Part of the IBM i Licensed Internal Code that manages the I/O operations to store and place data on storage
- Disk pools: Referred to as Auxiliary Storage Pools (ASP). These can be System ASPs, which contain required system information and without which IBM i will not start. Also, User ASPs are a way to separate storage for different objects in the system. Finally, Independent ASPs can switch among two or more IBM i systems that are in a cluster.

13.7.2 Fibre channel adapters and multipath

The DS8000 storage system can connect to the IBM i system in one of the following ways:

- Native: FC adapters in the IBM i system are connected through a Storage Area Network (SAN) to the Host Bus Adapters (HBAs) in the DS8000 storage system.

- With Virtual I/O Server Node Port ID Virtualization (VIOS NPIV): FC adapters in the VIOS are connected through a SAN to the HBAs in the DS8000 storage system.
- With VIOS: FC adapters in the VIOS are connected through a SAN to the HBAs in the DS8000 storage system.

Most installations use the native connection of the DS8000 storage system to the IBM i system or the connection with VIOS_NPIV.

Multipathing

IBM i multi-pathing is built into the IBM i System Licensed Internal Code (SLIC) and does not require a separate driver package to be installed. The IBM i system supports up to eight active connections (paths) to a single LUN in the DS8000 storage system.

When connecting the IBM i system to the DS8000 through VIOS, at least two VIOSs are required to implement Multipath. This provides additional resiliency if one VIOS fails. Also, the FC adapters in each VIOS can multipath to the connected DS8000 storage system to provide additional resiliency and enhance performance.

13.7.3 Performance guidelines for drives

The DS8000 supports both RAID 6 and RAID 10 configurations, with RAID 6 being the default and preferred option. However, due to a lower RAID penalty, RAID 10 may provide better performance in certain cases. Consider RAID 10 for IBM i systems especially for the following types of workloads:

- Workloads with large I/O rates
- Workloads with many write operations (low read/write ratio)
- Workloads with many random writes
- Workloads with low write-cache efficiency

13.7.4 Number and size of LUNs

For the performance of an IBM i system, ensure that the IBM i system uses many flash modules. The more flash modules that are available to an IBM i system, the more server tasks are available for the IBM i storage management to use for managing the I/O operations to the ASP space. The result is improved I/O operational performance.

It is important to provide many LUNs to an IBM i system, as the system treats these LUNs as physical disks or flash modules. Current versions of IBM i support up to 127 LUNs per FC adapter. Also, defining sufficiently small LUNs for an IBM i system is important, so as no to exceed the queue depth with I/O operations. Considering the manageability and limitations of external storage and an IBM i system, define LUN sizes of about 70.5 GB - 141 GB.

Extent pools for IBM i workload

Create two extent pools for an IBM i workload, with each pool in one rank group. The default extent allocation method (Rotate Extents) is recommended for IBM i systems, as it provides optimal performance compared to Rotate Volumes.

Note: Small extents are the recommended extent type for IBM i storage implementations.

Dedicating or sharing ranks

Sharing the ranks among the IBM i systems enables the efficient use of the DS8000 resources. However, the performance of each LPAR is influenced by the workloads in the

other LPARs. Deciding whether to share ranks to more efficiently utilize storage, dedicate ranks to IBM i LPAR to ensure stable performance, or a mix of the two depends on the nature of your workloads and environment. Be sure to use the proper planning tools, such as IBM StorM to determine the needs of your environment.

13.7.5 Analyzing performance data

The following tools are the IBM i performance data collectors. Collectors can be managed by IBM System Director Navigator for i, IBM System i Navigator, or IBM i commands.

Collection Services

Collection Services for IBM i can collect system and job level performance data. It can run all the time and will sample system and job level performance data with collection intervals as low as 15 seconds or alternatively up to an hour. It includes a set of specific statistics for the DS8000 storage system. For example, it shows which IBM i storage units are located within the DS8000 LUNs, whether they are connected in a single path or multipath, the disk service time, and wait time.

IBM i Job Watcher

Job Watcher is an advanced tool for collecting and analyzing performance information to help you effectively monitor your system or to analyze a performance issue. It is job-centric and thread-centric and can collect data at intervals of seconds.

IBM i Disk Watcher

Disk Watcher is a function of an IBM i system that provides disk data to help identify the source of disk-related performance problems on the IBM i platform.

Performance Explorer (PEX)

Performance Explorer is a data collection and analysis tool that helps identify performance issues that cannot be detected through Collection Services or general trend analysis. It enables detailed monitoring of application performance and isolates problems by pinpointing the specific system resource or application responsible for the issue.

Summary of best practices for IBM i host performance:

- Prefer RAID10 when workloads involve high random writes or low write-cache efficiency.
- Provide many LUNs and small LUN sizes. Define multiple LUNs (up to 127 per FC adapter) and keep sizes around 70–141 GB to avoid queue depth issues.
- Use at least two VIOS instances for multipath configurations and ensure multiple paths per LUN.
- Create two extent pools for IBM i workloads. Each pool should reside in a separate rank group for better performance and isolation.
- Prefer small extents for IBM i storage. Small extents improve capacity utilization and are recommended for Safeguarded Copy.
- Distribute workloads across multiple flash modules. More flash modules allow IBM i storage management to optimize I/O operations.
- Plan for proper rank sharing or isolation. Decide between sharing ranks for resource efficiency or dedicating ranks for predictable performance based on workload needs.
- Use IBM performance tools, such as Collection Services, Job Watcher, Disk Watcher, and Performance Explorer, to identify and resolve performance issues.

13.8 Host adapters

The DS8A00 models feature four-port 32 Gbps Fibre Channel/IBM FICON Host Adapters. These adapters comply with the Fibre Channel standard for backward compatibility, known as N-2 compatibility. As a result, a 32 Gbps adapter can automatically negotiate down to 16 Gbps or 8 Gbps. Host Adapters are available with either all shortwave (SW) or all longwave (LW) SFPs. Individual adapters do not support a mix of SW and LW ports; however, a DS8A00 system can include a mix of SW and LW adapters. Additionally, ports can be configured individually and dynamically, allowing FICON and FCP protocols to coexist on the same adapter.

The following lists the minimum and maximum host adapters by model:

- DS8A10
  - 2 - 16 (8-64 ports)
- DS8A50
  - 2 - 32 (8-128 ports)
- DS8A80
  - 2 - 32 (8-128 ports)

13.8.1 Fibre channel-attached hosts and SAN

Both direct connect and switched fabric topologies are supported in DS8A00 storage servers, with switch fabric providing maximum flexibility and performance. For specific considerations that apply to each server platform, and for the current information about supported servers (the list is updated periodically), refer to IBM System Storage Interoperation Center (SSIC).

Storage Attached Network (SAN)

A SAN shares centralized storage components, efficiently allocating storage to hosts via high-speed networking. Cabling hosts and network devices to different HAs in different I/O

bays when possible, as well as using multiple FC switches and directors ensures maximum availability and performance as well as avoids single points of failure.

Server restarts or hardware failures may cause an HA to repeatedly log into the SAN fabric with Registered State Change Notifications (RSCN). These RSCN have the potential to disrupt the I/O of other adapters. Therefore it is important to establish zones by either port number of WWPN to reduce the chance of interaction between the adapters. The preferred approach is a single initiator zone, which includes a single HA and multiple DS8A00 ports.

Important: A DS8000 HA port configured to run with FC topology may be configured to run in multiple zones, but it cannot be shared in a zone with z/OS hosts. Likewise, HA ports configured to run with FICON topology may not be shared in zones with non-z/OS hosts.

LUN masking

In FC attachment, LUN affinity is based on the WWPN of the adapter on the host. This LUN masking function on the DS8000 storage system is provided through the definition of DS8000 volume groups. By using the concept of LUN affinity coupled with limiting HA ports via FC zones, a host can see the same LUNs on the switched fabric via multiple paths.

SAN logical disk configuration

Physical paths: Each physical path to a logical disk on the DS8000 storage system is presented to the host operating system as a disk device.

In the SAN configuration shown in Figure 13-1 on page 483, the host connects to two SAN switches through two Fibre Channel (FC) adapters. Each SAN switch provides access, via zoning, to two DS8000 Host Adapter (HA) ports. This setup gives the host four distinct paths to every logical disk presented by the DS8000 storage system, meaning each logical disk appears to the host through four separate paths.

Careful planning of logical disk sizing and the number of paths from the host to the DS8000 storage server is essential to avoid an excessive increase in the number of logical devices presented to the host. For most small to mid-size servers, two to four paths per LUN provide an optimal balance of performance and manageability. In environments requiring very high data rates, up to eight paths may be considered. Adding more paths beyond this does not improve performance and only results in unnecessary complexity and an excessive number of devices being presented to the operating system.

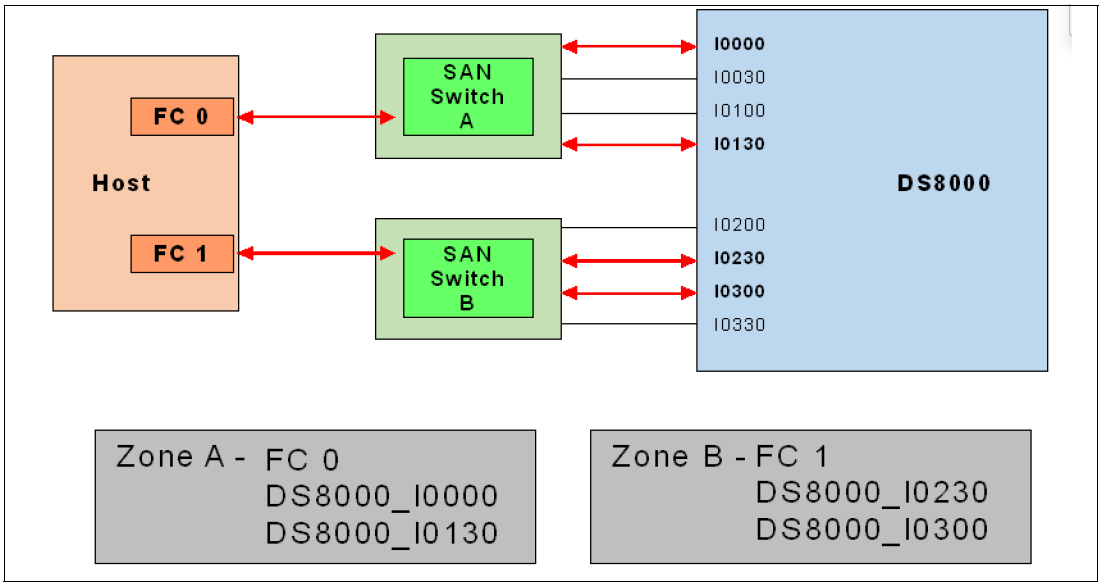


Figure 13-1 SAN zoning configuration

Multipathing

Multipathing describes a technique to attach one host to an external storage device through more than one path. Using multiple paths increases fault tolerance by avoiding single point of failure and increases system bandwidth, positively impacting system performance.

Multipathing software is necessary to properly utilize additional paths. Otherwise, the same LUN behind each path is handled as a separate disk from the operating system side, which does not allow failover support.

Multipathing provides the DS8000 attached Open Systems hosts with the following capabilities:

- Support for several paths per LUN
- Load Balancing between paths to avoid bottlenecks and improve performance
- Automatic path management, failover protection, and enhanced data availability
- Dynamic reconfiguration after environment changes

Multipath I/O (MPIO) summarizes native multipathing technologies that are available in several operating systems, such as AIX, Linux, and Windows. This is a module that is delivered with the OS and provides failover support, load balancing, and device-specific support. Implementations for each operating system may vary, but the basic concepts remain the same.

13.8.2 FICON and IBM Z hosts

FICON is a Fibre Connection used with IBM Z servers. FICON Host Adapters support 8, 16, or 32 Gbps link speeds in DS8A00 storage systems. When configured for FICON, the storage unit provides the following configurations:

- Either fabric or point-to-point topology.
- A maximum of 128 host ports for DS8A50 and DS8A80 models, and a maximum of 64 host ports for DS8A10 models.
- A maximum of 1280 logical paths per DS8000 HA port.
- Access to all 255 control unit images (up to 65,280 CKD devices) over each FICON port.

FICON connectivity supports IBM Z operating environments including z/OS, z/VM, Linux on IBM Z, VSE, and z/TPF.

Note: z/VM, VSE, and Linux on IBM Z can also connect to DS8000 storage systems using Fibre Channel Protocol (FCP). For these environments, the same best practices applied to Open Systems hosts, such as implementing multipathing for redundancy, ensuring high availability, and optimizing performance, should be followed.

FICON support:

- z15: FICON Express 16S, 16S+, 16SA
- z16: FICON Express 16S+, 16SA, 32S
- z17: FICON Express 16SA, 32S (2-port and 4-port cards)

Regardless of FICON generation, both direct-connect and switched fabric configurations are available for connecting DS8000 storage servers. See Figure 13-2 for examples of point-to-point, switched point-to-point, and cascaded fabric FICON topologies.

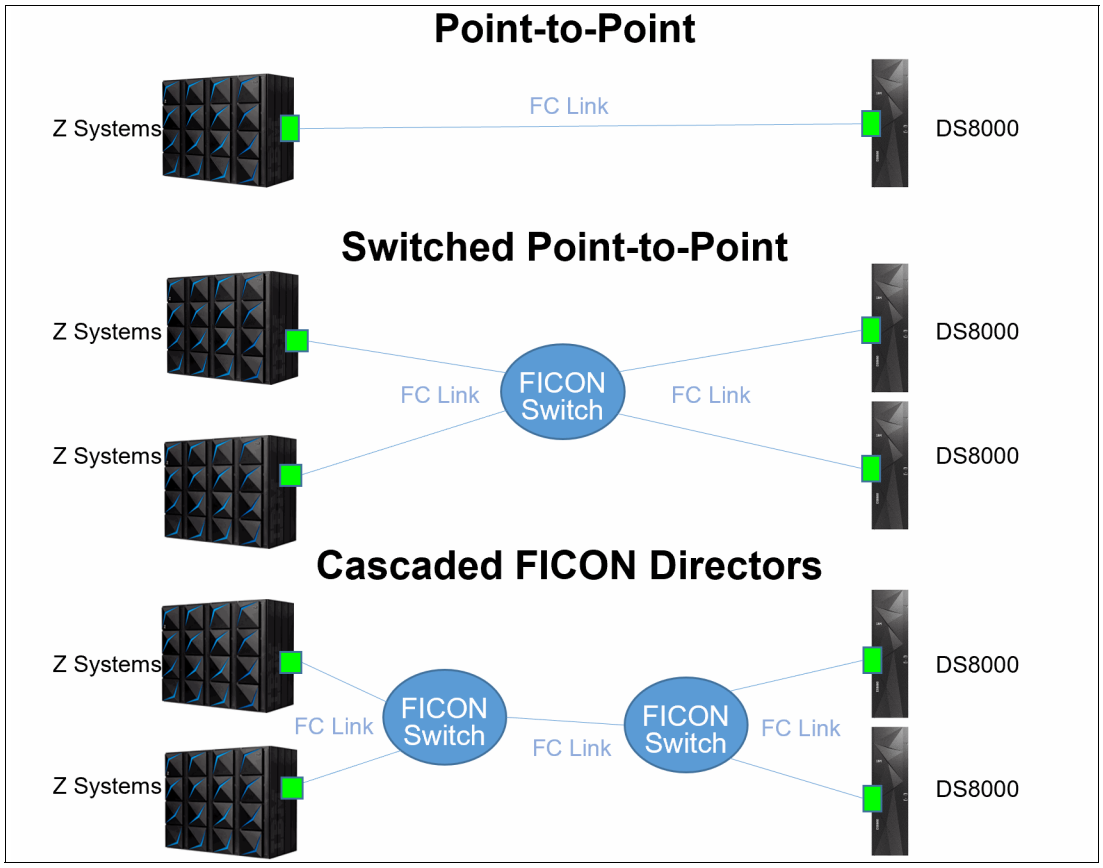


Figure 13-2 FICON topologies between Z Systems and a DS8000 storage system

Sizing

Sizing FICON connectivity can be challenging because it involves multiple factors. As a best practice, perform a detailed analysis of the specific environment before starting. Use the following guidelines when sizing the attachment environment:

- FICON Express CHPID utilization: Preferred maximum utilization level is 50%.
- FICON bus busy utilization: Preferred maximum utilization level is 40%.

- FICON Express link utilization: For estimated link speeds of 8 Gbps, 16 Gbps, or 32 Gbps, the preferred maximum utilization threshold is 70%.

It is also important to note that you should not mix paths with different link speeds to one LCU.

13.9 Practical performance management

There are a couple of major tools available to help you manage the performance of your DS8A00 storage server. These tools provide similar functionality but operate in slightly different ways.

- IBM Storage Insights:
  - IBM Cloud-based service with a data collector installed on-premises.
  - Free for use by IBM block storage customers, limited functionality
  - IBM Storage Insights Pro version to access full functionality
- IBM Spectrum Control:
  - Hosted on-premises
  - Does no send data to IBM
  - License includes access to IBM Storage Insights Pro at no additional charge

Note: For more information on IBM Storage Insights, refer to *Maximize Your Storage Investment: Use IBM Storage Insights for Observability and Artificial Intelligence for IT Operations*, SG24-8584.

Both IBM Spectrum Control and Storage Insights Pro monitor metrics for the following internal resources (see Table 13-5):

- Volumes
- Pools
- Arrays
- Nodes
- Fibre Channel Ports
- Host Connections

Table 13-5 IBM Spectrum Control and Storage Insights supported activities for performance processes

Process	Activities	Feature
Operational	Performance data collection for ports, arrays, volumes, pools, nodes (formerly called controllers), and host connections. Switch performance metrics can also be collected.	Performance monitor jobs.
Operational	Alerting.	Alerts and threshold violations.
Tactical/Strategic	Performance reporting of ports, pools, array, volumes, nodes, host connections, and switches.	Web-based GUI, IBM Cognos® (including predefined reports), and TPCTOOL.
Tactical	Performance analysis and tuning.	Tool facilitates thorough data collection and reporting.
Tactical	Short-term reporting.	GUI charting with the option to export data to analytical tools.

Process	Activities	Feature
Tactical	Advanced Analytics in IBM Spectrum Control.	Tiering and Balancing Analysis.
Strategic	Long-term reporting.	GUI charting with the option to export data to analytical tools.

14

Ransomware Threat Detection

This chapter provides an overview of Ransomware Threat Detection (RTD) capabilities for IBM DS8000 G10 storage systems. It explains how RTD leverages real-time analytics to identify ransomware-like behaviors, outlines the system requirements and licensing considerations, and describes the phased deployment approach. Additionally, it covers alert handling, integration with IBM Storage Insights Pro, and recommended response actions to strengthen cyber resiliency.

This chapter covers following topics:

- “Ransomware Threat Detection (RTD) overview” on page 488
- “RTD logical model” on page 489
- “Configure Threat Detection Storage (TDS)” on page 491
- “Storage Insights Pro for alerting” on page 494
- “Comparing IBM FlashSystem and DS8000: Key differences in Ransomware Threat Detection” on page 498

© Copyright IBM Corp. 2025.

14.1 Ransomware Threat Detection (RTD) overview

The IBM DS8000 G10 storage system, combined with the new FlashCore Module 4 (FCM4) technology, delivers advanced real-time telemetry and analytics at the storage layer. This capability continuously monitors I/O patterns and data entropy to identify ransomware-like behaviors, such as abnormal write spikes or irregular access patterns.

Note: For more information on FCM, refer to *IBM FlashCore Module (FCM) Product Guide: Features the newly available FCM4 with AI-powered ransomware detection*, REDP-5725.

By enabling the Real-Time Detection service, customers benefit from proactive protection through ongoing updates during the subscription period. In addition, the system provides immediate notifications when a potential threat is detected, allowing for rapid response and mitigation. This approach enhances data security by leveraging intelligent analytics directly within the storage infrastructure, reducing reliance on external monitoring tools.

Important: Ransomware threat detection requires the following components:

- FCM4 drives
- IBM DS8000 G10 storage system running code level 10.1 or higher
- Security Group License
- IBM Storage Insights Pro (for alerting in Phase 2)

Deployment phases:

- Phase 1 (Version 10.1): Data is transmitted for learning and model training purposes only. No alert notifications are generated for customers during this phase. Only IBM Support will see these alerts and they will be used for model tuning.
- Phase 2 (Version 10.2 and future releases): Customers with the Security Group License and IBM Storage Insights Pro will receive real-time alert notifications when potential threats are detected. The DS8000 GUI will not show ransomware alerts—the alerting mechanism will be through Storage Insights Pro and email notifications.

14.1.1 Key benefits

The following are the key benefits of ransomware threat detection:

- Proactive threat detection: Identifies ransomware-like behavior in real time using telemetry and entropy analysis.
- Immediate alerts: Provides instant notifications when anomalies or threats are detected.
- Integrated security: Embeds analytics at the storage layer, minimizing dependency on external solutions.
- Continuous updates: Ensures the latest detection capabilities throughout the subscription period.
- Reduced risk of data loss: Enables faster response and mitigation before significant damage occurs.
- Optimized performance: Maintains system efficiency while monitoring for threats without impacting workload performance.

14.1.2 Example scenario

A financial institution notices an unusual spike in write operations across multiple volumes during off-peak hours. The DS8000 G10 system, equipped with FCM4 and RTD service, detects this anomaly by analyzing I/O patterns and data entropy. The system immediately sends an alert to the storage administrator, indicating potential ransomware activity. The administrator can then take corrective actions, such as isolating affected volumes or initiating recovery procedures, before the malicious process encrypts critical data.

Immediate response workflow

The following high-level workflow outlines the key steps to isolate affected volumes, preserve a clean recovery point, validate data integrity, and restore operations—all with automation options to accelerate response time.

1. Isolate affected volumes:
   a. On IBM Z: Issue `VARY device(s),OFFLINE` or `VARY PATH` commands to remove the suspect volumes from active I/O.
   b. On distributed hosts: Unmap or mask the affected LUNs from host groups in DS8000, and optionally set them to read-only.
   c. Optional: Disable initiator WWPNs at the SAN fabric level for compromised hosts.
2. Capture an Immutable Recovery Point:
   a. Trigger an unplanned Safeguarded Copy backup via IBM Copy Services Manager to create an immutable snapshot of the current state.
   b. Optionally, use a FlashCopy Consistency Group with freeze for an additional point-in-time copy.
3. Examine in a Clean Room:
   a. Recover the safeguarded backup to dedicated recovery volumes using Copy Services Manager's Recover Backup function.
   b. Mount these recovery volumes on an isolated forensic host or LPAR for malware scans and data validation.
4. Restore production:
   a. After validation, use Copy Services Manager's Restore Backup to roll back production volumes to the last known-good safeguarded copy.
   b. Coordinate with remote copy services (Metro/Global Mirror, HyperSwap) before restore to avoid propagating corruption.
5. Automation considerations:
   a. Integrate Storage Insights Pro alerts with webhooks or SOAR platforms to automatically trigger CSM REST APIs for isolation and backup.
   b. Predefine Safeguarded Copy sessions and retention policies to ensure rapid execution.
   c. Maintain tested scripts for quick recovery and validation.

14.2 RTD logical model

The FCMs keeps track of I/O statistics on a monitor group that collects, aggregates and passes to the Inference Engine (AI Function) every 2 seconds to make decisions based upon training to determine if a ransomware threat is occurring on a monitor group basis.

FCM4 calculates entropy (estimate of randomness) and change in compression on every IOPS. It keeps statistics on each IOPS. FCM4 has two small RISC cores that process all this information. This information is statistically summarized into a relatively small amount of information per monitor group and the summaries are passed to an inference engine in DS8000.

Figure 14-1 shows the DS8000 Ransomware Detection logical model and how the entropy detection is then handled, firstly by the IBM DS8A00 code and then passed to IBM support if Call Home is configured and IBM Storage Insights Pro for alerting the customer.

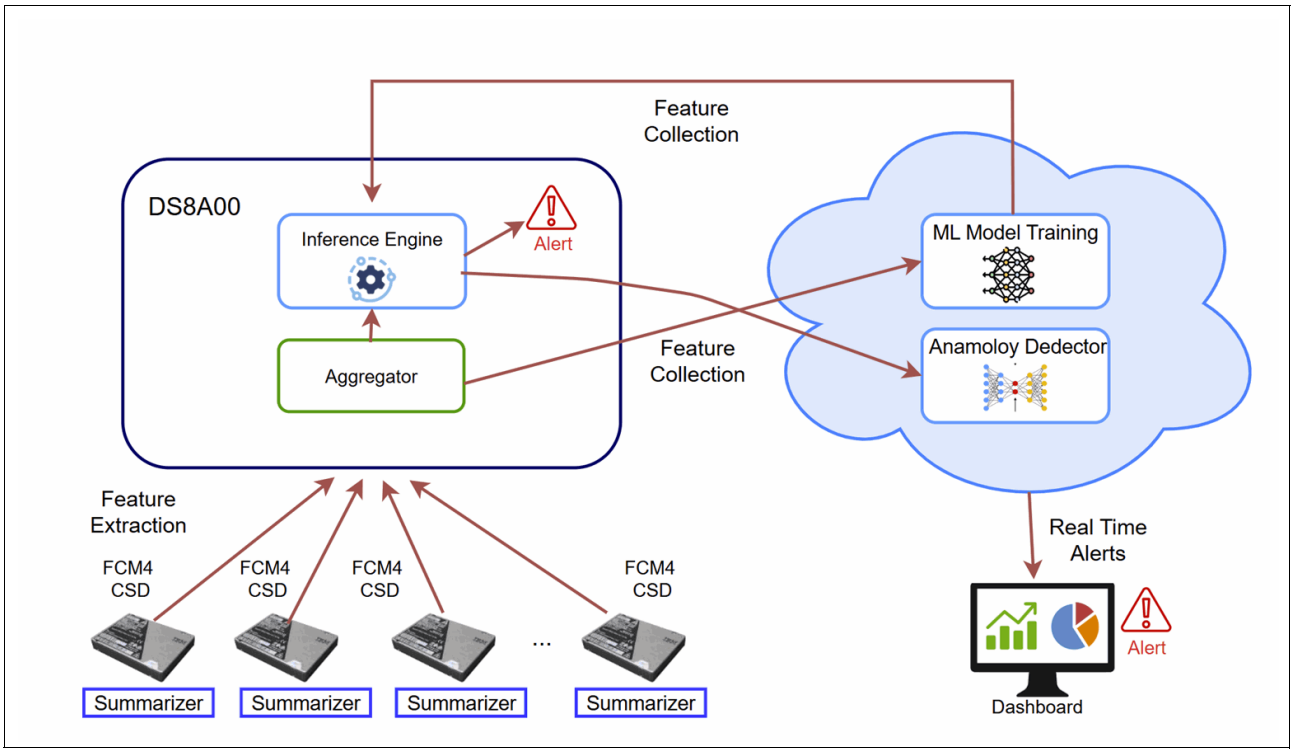


Figure 14-1 DS8000 Ransomware Detection conceptual model

Note: The Inference Engine monitors data from volumes mapped to a host. However, data from FlashCopy snapshots and Safeguarded Copy snapshots is not monitored. The reason for this is FlashCopy and Safeguarded Copy snapshots are designed to be read-only or point-in-time copies. Since they do not change after creation, monitoring them for ransomware-like behavior is unnecessary.

Here is the high-level workflow for implementing ransomware threat detection:

- **Data collection by Summarizers**
  - Each FCM4 collects per-write I/O operation metrics using a component called the Summarizer.
- **Feature aggregation**
  - Metrics from all FCM4 devices are sent to an Aggregator that resides within the DS8000 system's central processor complex (CPC).
  - The Aggregator combines these metrics into aggregated feature information.
- **Inference Engine analysis**

  - The Inference Engine runs on the DS8000 HMC
  - These engines will feed information from Machine Learning (ML) models on anomalies and ransomware attacks, it will learn what is normal for that system.
  - It will detect anomalous and dangerous behavior and take actions. This inference engine will be regularly updated.
- **ML model training**
  - New ML models are trained outside the storage system.
  - These models are periodically updated and retrained in IBM Storage Insights.
- **Feature collection and anomaly detection**
  - The Inference Engine applies anomaly detection using the trained ML models.
  - Alerts are generated when suspicious activity is detected.
- **Phased alerting**
  - Alerts are sent to the dashboard for real-time monitoring.
  - This enables administrators to respond to ransomware threats promptly.

14.3 Configure Threat Detection Storage (TDS)

Configure Threat Detection Storage components to manage data transmission and alert behavior. TDS functionality is managed through three configuration components that work together to define how data is transmitted and how alerts are generated.

14.3.1 Security Group License for IBM DS8000 systems

Security Group License affects customer notification capabilities:

- Purchased: In Phase 2, customers will receive real-time alert notifications when ransomware threats are detected.
- Not Purchased: Customers will not receive any alerts, even though data continues to be transmitted and analyzed for learning purposes.

14.3.2 Automatic Sending Threat Detection Data to IBM

This setting controls whether ransomware-related telemetry data is transmitted to IBM for analysis:

- Enabled (default setting): Ransomware detection data is automatically sent to IBM for advanced analytics and continuous improvement of detection models.
- Disabled: Data is not transmitted externally; however, local data generation continue within the DS8000 system.

Figure 14-2 illustrates the DS8000 Advanced Settings panel, where you can enable both Automatic Updates for Threat Detection Storage and Automatic Sending of Threat Detection Data to IBM.

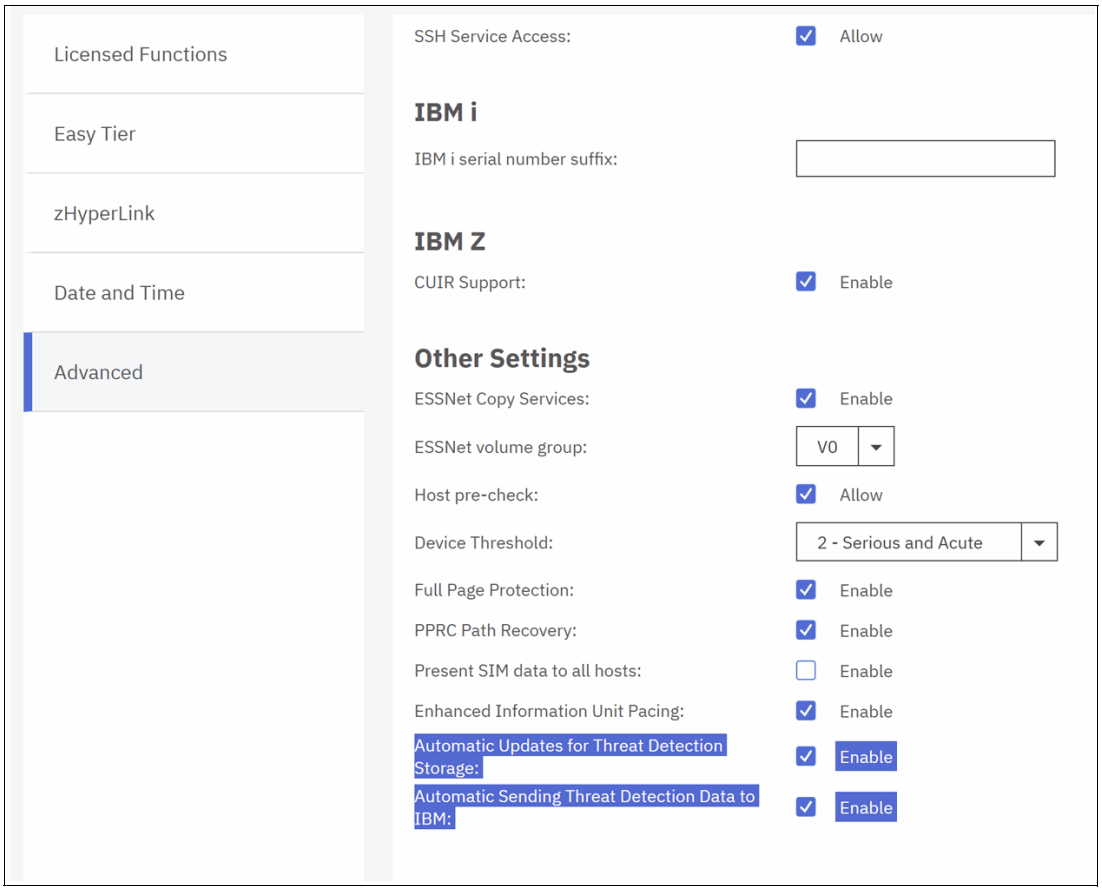


Figure 14-2 Settings to be enabled in DS8000

Notes:

- In DS8000 ransomware threat detection, alerts are generated by the Inference Engine within the system, but notifications are only sent through IBM Storage Insights. The DS8000 system itself does not issue local alert notifications to the customer—for example, when the ransomware data transmission switch is disabled. In such cases, data is still generated locally, but they are not transmitted or notified to the customer unless the Security LIC is installed and the transmission settings are enabled.
- Automatic Updates for Thread Detection Storage: This feature is enabled by default and ensures that the Inference Engine is updated automatically whenever necessary, keeping the detection logic current.

14.3.3 Call Home settings

The Call Home setting controls all data transmission between the DS8000 system and IBM.

- Enabled: Allows transmission of Call Home data, including ransomware-related data (if the Automatic Sending Threat Detection Data to IBM switch is also enabled).
- Disabled (default setting): Stops transmission of all Call Home data, including ransomware-related data and other information such as Advanced Call Home performance data.

Tip: By default, the “Automatic Sending Threat Detection Data to IBM” setting is enabled. This means the DS8000 system is configured to transmit ransomware-related telemetry data to IBM for advanced analytics and continuous improvement of detection models.

However, Call Home is disabled by default. Because Call Home governs all external data transmission, including ransomware telemetry, no data is actually sent to IBM unless Call Home is manually enabled by the administrator.

If data transmission is disabled, either by turning off the “Automatic Sending Threat Detection Data to IBM” setting or by leaving Call Home disabled, the system continues to generate ransomware detection data locally. These local folders are periodically cleaned to manage storage space and prevent unnecessary accumulation.

Figure 14-3 shows the DS8000 Call Home setting enabled for data transmission to IBM.

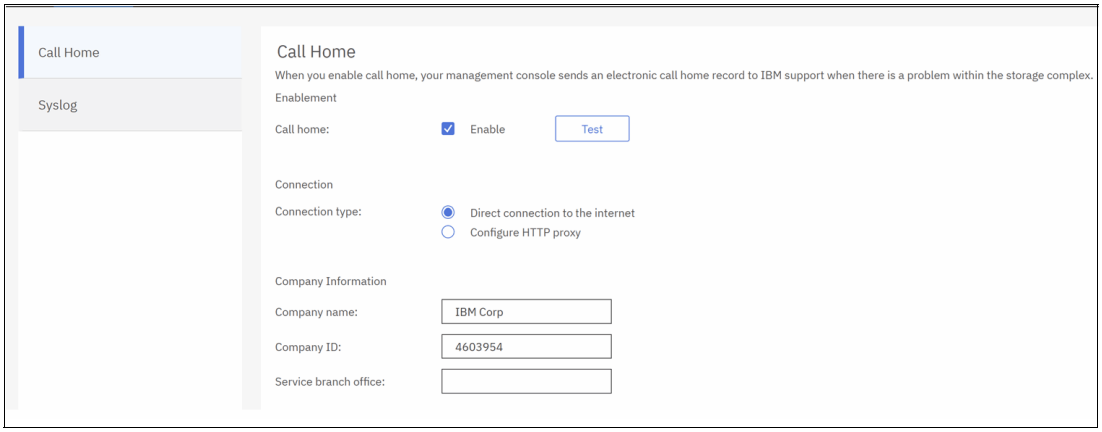


Figure 14-3 Call Home Settings

14.3.4 Threat Detection Storage reference

This reference provides configuration scenarios and transmission behavior for Threat Detection Storage components.

Figure 14-4 shows how different combinations of Security LIC, Automatic Sending Threat Detection Data to IBM Switch, and Call Home settings affect data transmission and customer notifications.

	Security LIC	Automatic Sending Threat Detection Data to IBM Switch	Call Home	Transmission	Notes
1	Installed	Enabled	Enabled	Yes ransomware data	Customer is notified of alerts
2	Installed	Disabled	Enabled	No ransomware data but Yes other Call Home data	Alerts and data are generated but never sent. Folders are cleaned periodically
3	Installed	Enabled	Disabled	No ransomware data No other call Home data	Alerts and data are generated but never sent. Folders are cleaned periodically
4	Installed	Disabled	Disabled	No transmission of any type	Alerts and data are generated but never sent. Folders are cleaned periodically
5	Not installed	Enabled	Enabled	Yes ransomware data	Customer is NOT notified of alerts. Alerts and data are generated but never sent. Folders are cleaned periodically
6	Not installed	Disabled	Enabled	No ransomware data but Yes other Call Home data	Customer is NOT notified of alerts. Alerts and data are generated but never sent. Folders are cleaned periodically
7	Not installed	Enabled	Disabled	No ransomware data No other call Home data	Customer is NOT notified of alerts. Alerts and data are generated but never sent. Folders are cleaned periodically
8	Not installed	Disabled	Disabled	No transmission of any type	Customer is NOT notified of alerts. Alerts and data are generated but never sent. Folders are cleaned periodically

Figure 14-4 Threat Detection Storage reference

14.4 Storage Insights Pro for alerting

IBM Storage Insights offers a modern, analytics-driven, and centralized approach to storage monitoring and optimization. This cloud-based service enhances visibility, reduces maintenance time, and optimizes storage performance through AI-powered recommendations. It also helps detect and respond to potential ransomware attacks by forwarding alerts directly to you, enabling rapid action to minimize potential damage.

Important: Full ransomware alerting capabilities are available in Phase 2 of IBM's Threat Detection rollout and require IBM Storage Insights Pro. The free version of Storage Insights does not support ransomware threat detection features.

IBM Storage Insights Pro is included with Expert Care Premium for the duration of the Expert Care entitlement. You do not need to purchase it separately if you have Expert Care Premium. Expert Care Basic or Advanced Includes the free version of IBM Storage Insights only.

Other entitlement options for IBM Storage Insights Pro include IBM Spectrum Control, IBM Storage as a Service (STaaS), or separate subscription purchase.

For more information on DS8A00 integration with Storage Insights, refer to *Maximize Your Storage Investment: Use IBM Storage Insights for Observability and Artificial Intelligence for IT Operations*, SG24-8584.

14.4.1 Ransomware alert details

Note: This section is available only to IBM Support team in the initial release (Phase 1).

Sample ransomware alert displayed in the IBM Storage Insights console. See Figure Figure 14-5 on page 495.



Figure 14-5 1 IBM Storage Insights alerts

The alert displayed in Figure 14-6 on page 496 includes the following details:

- Severity: Indicates that ransomware threats have been detected.
- Alert Source: The name of the storage system where the threat was identified.
- Instructions:
  - Steps to validate a ransomware alert.
  - Actions to take if the alert is confirmed as ransomware.
  - Procedures to follow if the alert is determined to be a false positive.

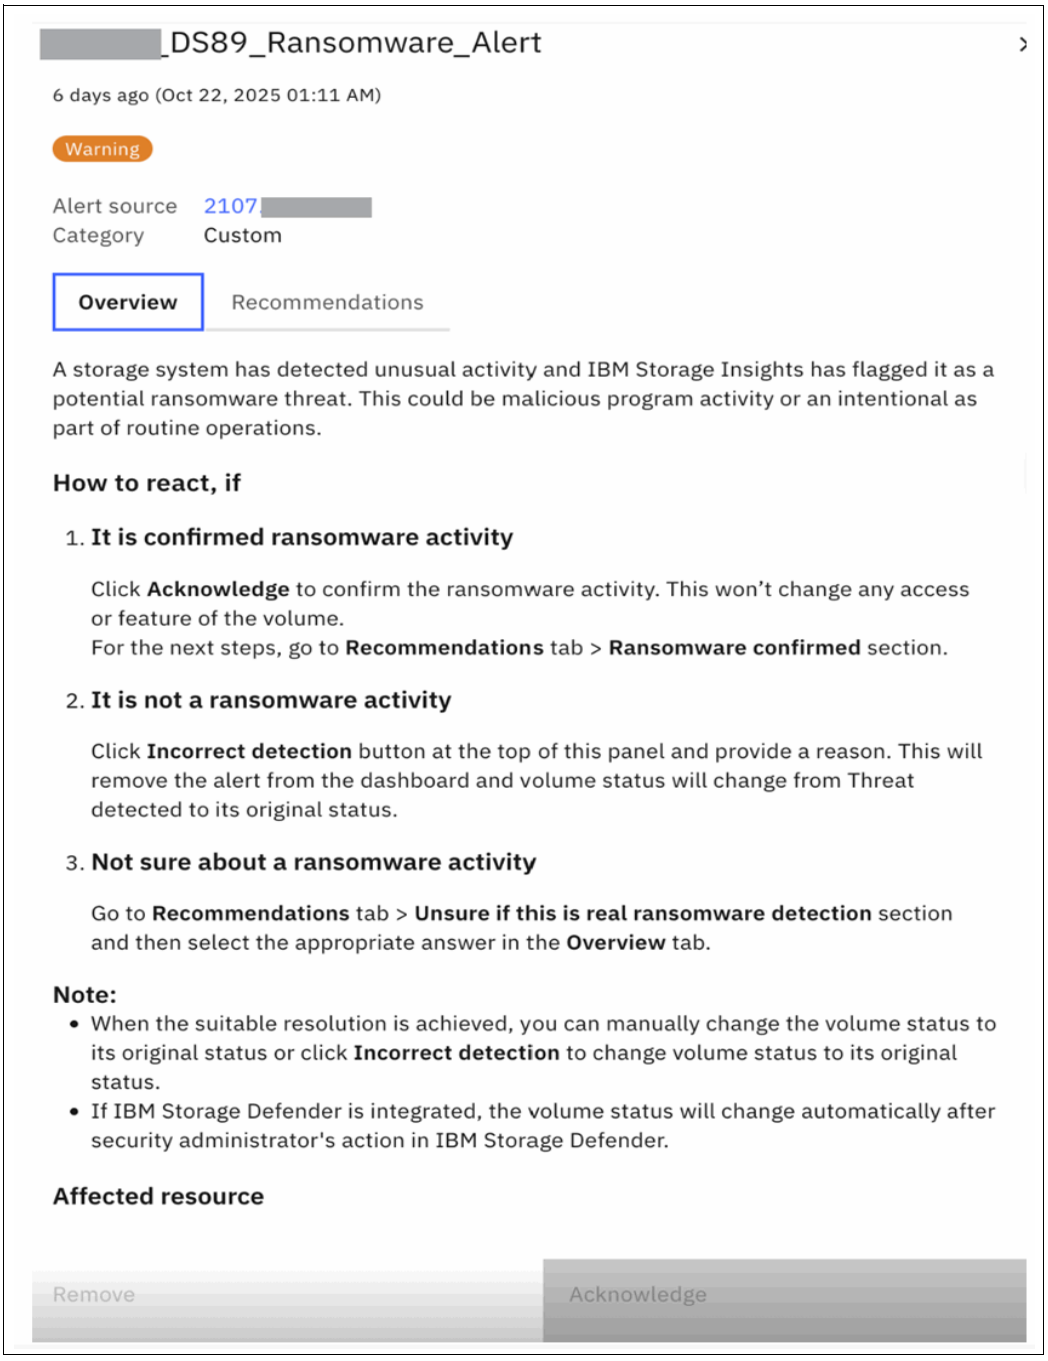


Figure 14-6 IBM Storage Insights ransomware threat overview

When a ransomware threat detection alert is triggered, the Alert dialog box provides best practices (Recommendations), as shown in Figure 14-7 on page 497.

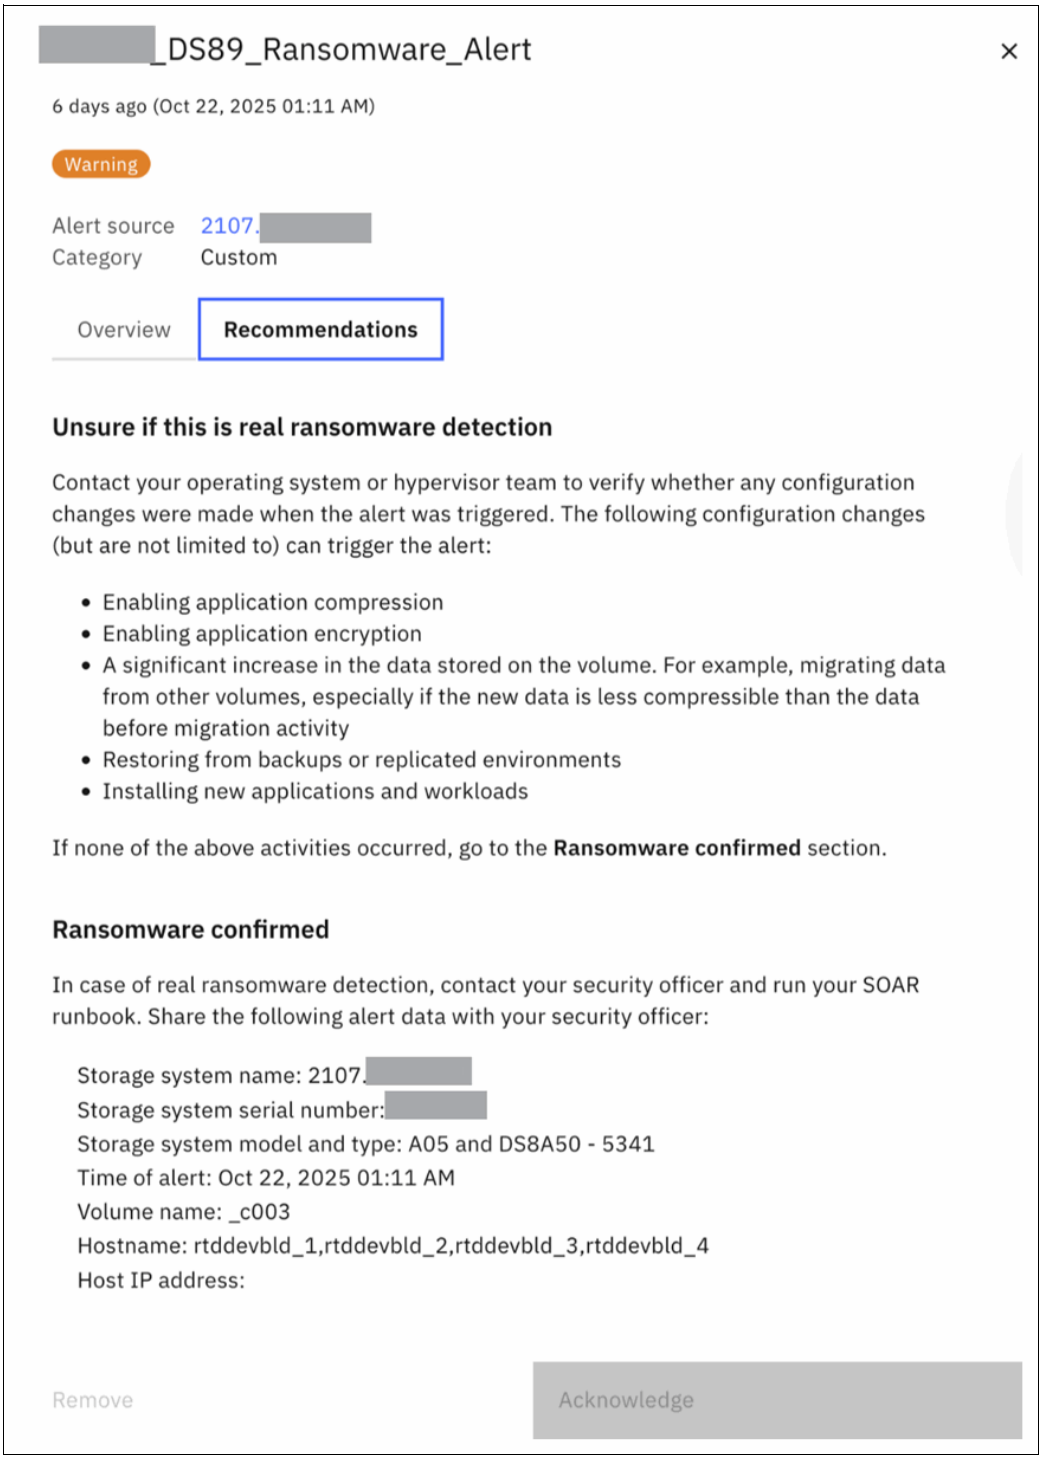


Figure 14-7 IBM Storage Insights ransomware threat Recommendations

For example, the administrator should first validate whether recent configuration changes are the cause of the alert. This validation might involve contacting the operating system administrator or the hypervisor team. Possible causes for the alert (including but not limited to) might include the following items:

- Enabling application compression, such as compressing many files.

- Enabling application encryption, particularly for application files containing sensitive data, such as financial information.
- An increase in stored data on the volume, such as data migration from other volumes, especially if the new data is less compressible than the previous data.
- Restoring from backups or replicated environments.
- Installing new applications and workloads.

If none of the activities that are described in Figure 14-7 on page 497 has occurred and a genuine ransomware threat is detected, contact your security officer and provide them with the details in Figure 14-7.

This information includes the storage system name, model, and serial number, and recommends running the security, orchestration, automation, and response (SOAR) runbook.

The SOAR runbook might include activities such as (but not limited to) the following ones:

- Initiating automated data collection.
- Helping ensure the continuation of organizational mission and business functions.
- Automatically disabling impacted systems.

Note: Users can configure Security SOAR tools by using the SI webhook configuration. See Creating a webhook in IBM Storage Insights for more information.

For more information, see Security and Privacy Controls for Information Systems and Organizations, NIST Special Publication 800-53.

14.5 Comparing IBM FlashSystem and DS8000: Key differences in Ransomware Threat Detection

IBM FlashSystem and DS8000 both provide ransomware protection, but their implementations differ in architecture, deployment, and operational behavior. These variations influence licensing requirements, integration with security workflows, and recovery strategies. Table 14-1 gives a summary of these differences.

Table 14-1 IBM FlashSystem versus DS8000: How their Ransomware Detection approaches differ

Feature	FlashSystem	DS8000
Detection location	Inline at data path (FCM4 drives)	Storage layer via TDS service
License requirement	None for detection and alerts	Security LIC required for alerts
Integration	IBM Storage Insights and IBM Storage Defender	IBM Storage Insights and SOAR runbooks
Granularity	Block-level	System-level
Deployment	Fully active after registration	(*) Two-phase rollout
Recovery features	(**) Uses Safeguarded Copy to create immutable snapshots (up to 256 immutable backups per volume).	(***) Also uses Safeguarded Copy, (up to 1024 immutable backups per volume).

(*) Phase 1 – Learning Mode (Code Level 10.1)

- The system collects and transmits telemetry data to IBM for model training and behavior analysis.

- No alerts are sent to customers during this phase.

- Purpose: Build accurate detection models without generating false positives.

Phase 2 – Alerting Mode (Code Level 10.2 and later)

- Once the detection models are mature, the system begins real-time alerting.

- Customers with the Security Group License (Security LIC) receive notifications when ransomware-like activity is detected.

- Enables proactive response and integration with SOAR workflows.

(**) FlashSystem implementation uses Safeguarded Copy to create immutable snapshots that cannot be altered or deleted by host or user actions, including ransomware attacks. These safeguarded snapshots are scheduled via internal policies or external tools like IBM Copy Services Manager.

(***) DS8000 implementation also uses Safeguarded Copy, but the implementation differs in scale and architecture (up to 1024 immutable backups per volume). DS8000 does not have an internal scheduler for Safeguarded Copy. All scheduling and management of safeguarded backups must be done through IBM Copy Services Manager, which is pre-installed on the DS8000 HMC, so no separate server is needed.

Abbreviations and acronyms

AC	active
ACK	acknowledgment
AIXPCM	AIX Multipath I/O Path Control Module
AMP	Adaptive Multi-stream Prefetching
ANSI	American National Standards Institute
AOS	Assist On-site
API	application programming interface
ASIC	application-specific integrated circuit
ATS	Atomic Test and Set
BF	Base Function
BOS	base operating system
BPM	Backup Power Module
BTU	British thermal unit
CA	certificate authority
CCL	concurrent code load
CDA	Code Distribution and Activation
CEC	Central Electronics Complex
CIM	Common Information Model
CKD	Count Key Data
CMA	cable management arm
CPC	central processor complex
CRC	cyclic redundancy check
CS	Copy Services
CSI	Container Storage Interface
CSR	certificate signing request
CSV	comma-separated values
CUADD	control unit address
CUIR	control-unit initiated reconfiguration
DA	device adapter
DC-UPS	direct current uninterruptible power supply
DDM	disk drive module
DIF	Data Integrity Field
DIMM	dual inline memory module
DMA	direct memory access
DPR	Dynamic Path Reconnect
DPS	Dynamic Path Selection
DR	disaster recovery
DRP	disaster recovery planning
DSFA	data storage feature activation
DSNI	DS Network Interface
DVE	Dynamic Volume Expansion
DVR	Dynamic Volume Relocation
DWDM	dense wavelength division multiplexing
EAM	extent allocation method
EAV	extended address volume
ECC	error correction code
eDRAM	embedded dynamic random access memory
EGP	Exterior Gateway Protocol
EKM	Encryption Key Manager
EOS	end of service
EPOW	emergency power-off warning
ER	Energy Report
ESE	extent space efficient
ESM	Enclosure Service Module
FB	Fixed-Block
FC	Fibre Channel
FC-AL	Fibre Channel Arbitrated Loop
FCIC	Fibre Channel Interface Card
FCP	Fibre Channel Protocol
FDE	Full Disk Encryption
FEC	Forward Error Correction
FFDC	first-failure data capture
FICON	Fibre Channel connection
FIDR	FICON Dynamic Routing
FPCCL	fast path concurrent code load
FQDN	fully qualified domain name
FRU	field-replaceable unit
FSP	flexible service processor
GC	Global Copy
GFC	gigabit Fibre Channel
GiB	gibibyte
GM	Global Mirror
HA	high availability
HBA	host bus adapter
HCD	hardware configuration definition
HDD	hard disk drive

© Copyright IBM Corp. 2025.

HMC	Hardware Management Console
HMT	Heat Map Transfer
HMTU	Heat Map Transfer Utility
HPFE	High-Performance Flash Enclosure
I2C	Inter-Integrated Circuit
IBM	International Business Machines Corporation
IBM ESSNI	Enterprise Storage Server Network Interface
IBM SSR	IBM Systems Service Representative
ICS	Install Corrective Service
IML	initial microcode load
IOCDS	input/output configuration data set
IODF	input/output definition file
IOPS	I/O operations per second
iPDU	intelligent Power Distribution Unit
IPSec	Internet Protocol Security
IRR	instruction retry and recovery
ISL	inter-switch link
IWC	Intelligent Write Caching
JCL	Job Control Language
JVM	Java virtual machine
KMIP	Key Management Interoperability Protocol
kW	kilowatt
LAN	local area network
LBA	logical block address
LCP	Logical Corruption Protection
LCU	logical control unit
LD	logical device
LDAP	Lightweight Directory Access Protocol
LIC	Licensed Internal Code
LMC	Licensed Machine Code
LPAR	logical partition
LSS	logical subsystem
LUN	logical unit number
LW	longwave
MBps	megabytes per second
MC	Management Console
ME	Management Enclosure
MGM	Metro/Global Mirror
MIB	Management Information Base
MiB	mebibyte
MIPS	million instructions per second
MM	Metro Mirror
MPIO	multipath I/O
MRPD	machine-reported product data
MSDSM	Microsoft Device Specific Module
MTM	machine type and model
MTMM	Multiple-target Metro Mirror
NMP	Native Multipathing Plug-in
NTP	Network Time Protocol
NVDIMM	non-volatile dual inline memory module
NVS	non-volatile storage
OCP	Red Hat OpenShift Container Platform
ODD	OnDemand Data Dump
OEL	Operating Environment License
OID	object identifier
OS	operating system
PAV	parallel access volume
PCIe	Peripheral Component Interconnect Express
PCN	Power Control Network
PFA	predictive failure analysis
PFE	Product Field Engineer
PHYP	Power9 IBM PowerVM Hypervisor
PKCS	Public-Key Cryptography Standard
PLD	power line disturbance
POST	power-on self-test
PPRC	Peer-to-Peer Remote Copy
PSU	power supply unit
PTC	point-in-time copy
PTF	Program Temporary Fix
QoS	quality of service
RAID	redundant array of independent disks
RAS	reliability, availability, and serviceability
RC	reference code
RCD	register clock driver
RCL	Remote Code Load
RDP	Read Diagnostic Parameters
REST	Representational State Transfer
RMZ	Remote Mirror for z/OS
RPC	Remote Procedure Call
RPCC	rack power control card

RPM	revolutions per minute
RPQ	Request for Price Quotation
RSC	Remote Support Center
RSCN	registered state change notification
SAM	storage allocation method
SAN	storage area network
SARC	Sequential Adaptive Replacement Cache
SCM	single-chip module
SCORE	Storage Customer Opportunity Request
SCSI	Small Computer System Interface
SDD	Subsystem Device Driver
SDDDSM	Subsystem Device Driver Device Specific Module
SDDPCM	Subsystem Device Driver Path Control Module
SDM	System Data Mover
SDO	Secure Data Overwrite
SDS	software-defined storage
SES	Small Computer System Interface Enclosure Services
SFF	small form factor
SFI	storage facility image
SFP	small form-factor pluggable
SFP+	small form-factor pluggable plus
SIM	service information message
SLB	segment lookaside buffer
SMP	symmetric multiprocessor
SMT	simultaneous multithreading
SNMP	Simple Network Management Protocol
SOI	Silicon-On-Insulator
SPOF	single point of failure
SR	suspension reason
SRA	Storage Replication Adapter
SRC	system reference code
SSD	solid-state drive
SSH	Secure Shell
SSIC	System Storage Interoperation Center
SSID	subsystem identifier
SSL	Secure Sockets Layer
SSO	single sign-on
SW	shortwave
TCE	Translation Control Entry
TCO	total cost of ownership
TCT	Transparent Cloud Tiering
TLS	Transport Layer Security
TSO	Time Sharing Option
UA	unit address
UCB	unit control block
UPS	uninterruptible power supply
URI	Unified Resource Identifier
VAAI	vStorage APIs for Array Integration
VM	virtual machine
VPN	virtual private network
WLM	Workload Manager
WUI	Web User Interface
WWNN	worldwide node name
WWPN	worldwide port name
zDDB	z/OS Distributed Data Backup
zGM	z/OS Global Mirror
zHL	zHyperLink
zHPF	High-Performance FICON for IBM Z
zsS	Z Synergy Services

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 Cascading FlashCopy Design and Scenarios*, REDP-5463
- *DS8000 Global Mirror Best Practices*, REDP-5246
- *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*, SG24-8367
- *IBM DS8000 Easy Tier*, REDP-4667
- *IBM DS8000 Encryption for Data at Rest, Transparent Cloud Tiering, and Endpoint Security*, REDP-4500
- *IBM DS8000 and IBM Z Synergy*, REDP-5186
- *IBM DS8000 Safeguarded Copy*, REDP-5506
- *IBM DS8000 and Transparent Cloud Tiering*, SG24-8381
- *IBM DS8870 and NIST SP 800-131a Compliance*, REDP-5069
- *IBM DS8000 Thin Provisioning*, REDP-5343
- *IBM DS8900F Performance Best Practices and Monitoring*, SG24-8501
- *IBM DS8900F Product Guide*, REDP-5554
- *IBM DS8910F Model 993 Rack-Mounted Storage System*, REDP-5566
- *IBM Fibre Channel Endpoint Security for IBM DS8900F and IBM Z*, SG24-8455
- *IBM FlashCore Module (FCM) Product Guide*, REDP-5725
- *IBM Power Systems S922, S914, and S924 Technical Overview and Introduction Featuring PCIe Gen 4 Technology,* REDP-5595
- *IBM TS7700 Series DS8000 Object Store User's Guide*, REDP-5583
- *IBM Storage DS8900F Architecture and Implementation*, SG24-8456
- *IBM System Storage DS8000: Remote Pair FlashCopy (Preserve Mirror)*, REDP-4504
- *IBM System Storage DS8000: z/OS Distributed Data Backup*, REDP-4701

© Copyright IBM Corp. 2025. All rights reserved.

- *LDAP Authentication for IBM Storage DS8000 Systems*, REDP-5460
- *Maximize Your Storage Investment: Use IBM Storage Insights for Observability and Artificial Intelligence for IT Operations*, SG24-8584
- *IBM B-Type SAN Extension Platform Implementation Guide*, SG24-8553
- *IBM Storage Networking c-type FICON Implementation Guide*, SG24-8468
- *IBM Z Connectivity Handbook*, SG24-5444
- *IBM z17 Technical Introduction*, SG24-8580
- *IBM z17 (9175) Technical Guide*, SG24-8579
- *IBM z16 (3931) Technical Guide*, SG24-8951
- *IBM z16 A02 and IBM z16 AGZ Technical Guide*, SG24-8952
- *IBM z17 Configuration Setup*, SG24-8581

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 Storage DS8000 10.1 Introduction and Planning Guide, G10-I53-01
  https://www.ibm.com/docs/en/SSY1BJD_10.1.0/g10_i53_01_20250825.pdf
- *IBM DS8000 Series Command-Line Interface User's Guide*, G10-C53-01
  https://www.ibm.com/docs/en/SSY1BJD_10.1.0/g10_c53_01_20251007.pdf
- *IBM Storage DS8000 10.0 Host Systems Attachment Guide*, SC27-9563
  https://www.ibm.com/docs/en/SSY1BJD_10.0/sc27956302.pdf
- *RESTful API Guide, SC27-9823*
  https://publibfp.dhe.ibm.com/epubs/pdf/c2798230.pdf

Online resources

These websites are also relevant as further information sources:

- DS8000 IBM Documentation:

  https://www.ibm.com/docs/en/ds8000/10.1.0
- IBM Entitled Systems Support (ESS):

  https://www.ibm.com/servers/eserver/ess
- DS8000 Series Copy Services Fibre Channel Extension Support Matrix:

  https://www.ibm.com/support/pages/ds8000-series-copy-services-fibre-channel-extension-support-matrix
- IBM Fix Central:

  https://www.ibm.com/support/fixcentral/swg/selectFixes?parent=Enterprise%20Storage%20Servers&product=ibm/Storage_Disk/DS8A00&release=All&platform=All&function=all

- IBM System Storage Interoperation Center (SSIC) for DS8000:
  https://www.ibm.com/systems/support/storage/ssic/
- IBM DS8000 Code Recommendations:
  https://www.ibm.com/support/pages/ds8000-code-recommendation
- DS8000 Code Bundle Information (includes Release Notes):
  https://www.ibm.com/support/pages/node/7174030
- DS8000 Host Adapter Configuration Guidelines:
  https://www.ibm.com/support/pages/ds8000-host-adapter-configuration-guidelines
- DS8000 Service R10.*x* (available via IBM-ID login):
  https://ibmdocs-test.dcs.ibm.com/docs/en/ds8000-s/10.1.0
- IBM Remote Code Load:
  https://www.ibm.com/support/pages/ibm-remote-code-load

Help from IBM

IBM Support and downloads

ibm.com/support

Services from IBM Consulting

ibm.com/services

IBM Training and Skills

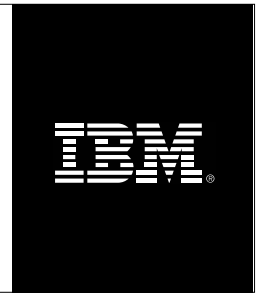
ibm.com/training



IBM Storage DS8A00 Architecture and

SG24-8559-01

ISBN



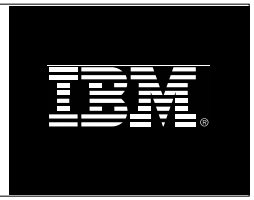
(1.5” spine)
1.5”<-> 1.998”
789 <->1051 pages



IBM Storage DS8A00 Architecture and Implementation

SG24-8559-01

ISBN



(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

Redbooks

IBM Storage DS8A00 Architecture and Implementation Guide:

SG24-8559-01

ISBN



(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

Redbooks

IBM Storage DS8A00 Architecture and Implementation Guide: Updated for

(0.2”spine)
0.17”<->0.473”
90<->249 pages

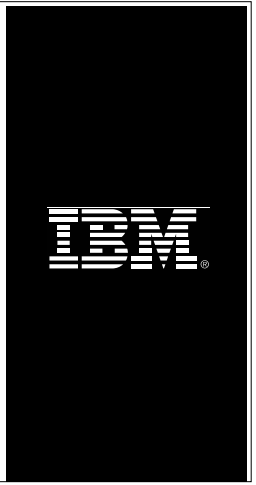
(0.1”spine)
0.1”<->0.169”
53<->89 pages



IBM Storage DS8A00 Architecture and

SG24-8559-01

ISBN



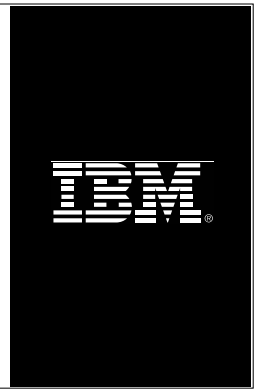
(2.5” spine)
2.5”<->nnn.n”
1315<-> nnnn pages



IBM Storage DS8A00 Architecture and Implementation

SG24-8559-01

ISBN



(2.0” spine)
2.0” <-> 2.498”
1052 <-> 1314 pages



SG24-8559-01

ISBN

Printed in U.S.A.

Get connected



